



Installing Cray System Management Workstation (SMW) Software, S-2480-UP03

Contents

Introduction.....	4
Related Publications.....	4
Distribution Media.....	5
Install the SMW Release Package.....	6
Before You Begin.....	6
Network Connections.....	6
Configuration Values.....	7
Passwords.....	8
Back Up Your Current Software.....	9
Rack-mount SMW: Install the base operating system.....	9
Desk-side SMW: Install the base operating system.....	15
Install and Configure the SMW Software Packages.....	18
Shut down the Cray system.....	18
Install and configure SMW software	18
Discover the Cray System Hardware and Power Up the Full System.....	21
Bootstrap Hardware Discovery.....	21
Discover the Cray System Hardware.....	25
Confirm the SMW is Communicating with System Hardware.....	30
Change the Default SMW Passwords.....	32
Set up the SUSE firewall and IP tables.....	32
Back Up Your Newly-installed SMW Software.....	33
Rack-mount SMW: Create an SMW bootable backup drive.....	33
Rack-mount SMW: Set Up the Bootable Backup Drive as an Alternate Boot Device.....	41
Desk-side SMW: Create an SMW Bootable Backup Drive.....	43
Desk-side SMW: Set Up the Bootable Backup Drive as an Alternate Boot Device.....	50
Configure the Simple Event Correlator (SEC).....	51
Configure Boot RAID.....	53
Prerequisites and Assumptions for Configuring the Boot RAID	53
Configure and Zone the Boot RAID Using a DDN Storage System.....	53
Configure a DDN Storage System for the Boot RAID via telnet.....	54
Install the Silicon Storage Appliance Manager GUI	58
Configure Remote Logging of DDN Messages.....	61
Configure the Boot RAID Using a NetApp, Inc. Engenio Storage System.....	62
Install SANtricity Storage Manager Software for NetApp, Inc. Engenio Devices.....	62
Configure the LUNs for NetApp, Inc. Engenio Devices.....	64

Configure Remote Logging of NetApp, Inc. Engenio Storage System Boot RAID Messages.....	66
Zone the QLogic™ FC Switch.....	66
Configure zoning for a QLogic SANbox switch using QuickTools utility.....	67
Create a backup of your QLogic switch configuration.....	68
Rediscover the LUNs.....	68
Reboot the SMW and verify LUNs are recognized.....	68
Partition the LUNs.....	69
Set Up Boot RAID Space for Direct-attached Lustre.....	69
Create LUNs and an IMPS file system for DAL.....	69
Update or Upgrade Your Cray SMW Software.....	71
Before You Start.....	71
Back Up Your Current Software.....	72
Rack-mount SMW: Create an SMW bootable backup drive.....	72
Desk-side SMW: Create an SMW Bootable Backup Drive.....	80
Install an SMW Update Package.....	86
Shut down the Cray system.....	87
Update the SMW software and configuration.....	87
For Cray XC Series Systems Only: Update the BC and CC Firmware.....	89
For Cray XE or Cray XK Systems Only: Verify Hardware Programmable Interface Controllers (PICs).....	92
For Cray Systems with Cray XK Blades Only: Update the Cray XK Accelerator Firmware.....	93
Update SMW Software on the Boot Root and Shared Root.....	93
Install Additional Software as RPMs.....	96
Enable Remote Access to the SMW.....	97
Start the VNC server.....	97
For workstation or laptop running Linux: Connecting to the VNC server through an SSH tunnel, using the vncviewer -via option.....	98
For workstation or laptop running Linux: Connecting to the VNC server through an SSH tunnel.....	99
For workstation or laptop running Mac OS X: Connecting to the VNC server through an SSH tunnel.....	99
For workstation or laptop running Windows: Connecting to the VNC server through an SSH tunnel.....	100
Set up the SUSE firewall and IP tables.....	101
Enable an Integrated Dell Remote Access Controller (iDRAC6) on a rack-mount SMW.....	103
Change the default iDRAC Password.....	105
Using the iDRAC6.....	106
Rack-mount SMW: Install the Dell Systems Management Tools and Documentation DVD.....	107
Rack-mount SMW: Configure the SMW to restore AC power after a power failure.....	108
Rack-mount SMW: Replace a failed disk drive.....	109

Introduction

This guide contains procedures for installing Cray System Management Workstation (SMW) operating software on desk-side and rack-mount SMW hardware. Procedures are provided for:

- Performing an initial installation of the SMW release package. The SMW software consists of the SMW base operating system, which is SUSE® Linux Enterprise Server version 11 Service Pack 3 (SLES® 11 SP3), and the Cray SMW software. See [Install the SMW Release Package](#) on page 6.
- Configuring the boot RAID; see [Configure Boot RAID](#) on page 53.
- Updating your Cray SMW software. See [Update or Upgrade Your Cray SMW Software](#) on page 71.

IMPORTANT: Before proceeding with any upgrade or update installation read the release notes to determine supported upgrade paths from previous versions of the SMW.

The intended user of this guide is a system administrator who is familiar with operating systems derived from UNIX.

Related Publications

Although this document is all that is necessary for the SMW software installation, the following documents contain additional information that may be helpful:

- *SMW Release Errata* and the *SMW README*, which are provided with your SMW release package
- *Cray Linux Environment (CLE) Software Release Overview (S-2425)*, which is provided with your CLE release package
- *Cray Linux Environment (CLE) Software Release Overview Supplement (S-2497)*, which is provided with your CLE release package as of its own update release package
- *Installing and Configuring Cray Linux Environment (CLE) Software (S-2444)*, which is provided with your CLE release package
- *Managing System Software for the Cray Linux Environment (S-2393)*, which is provided with your CLE release package
- *Managing Lustre for the Cray Linux Environment (CLE) (S-0010)*, which is provided with your Cray Linux Environment (CLE) operating system release package
- *Using and Configuring System Environment Data Collections (SEDC) (S-2491)*, which is provided with your SMW release package
- *Cray Programming Environments Installation Guide (S-2372)*, which is provided with your Cray Developer's Toolkit (CDT) release package
- *Monitoring and Managing Power Consumption on the Cray XC30 System (S-0043)*
- *Network Resiliency for Cray XC30 Systems (S-0041)*

Distribution Media

The Cray SMW 7.2 release distribution media consists of two DVDs.

- If you are performing an initial installation of the SMW software, the first DVD you will install is labeled `Cray-SMWbase11SP3-yyyymmdd` and contains the SMW base operating system, SUSE Linux Enterprise Server version 11 Service Pack 3 (SLES 11 SP3). The second is labeled `Cray SMW 7.2UPxx` and contains the Cray SMW 7.2 software packages.
- If you are updating your SMW software, use only the DVD labeled `Cray SMW 7.2UPxx`, which contains the Cray SMW 7.2 software packages. All software is installed by means of scripts and RPM Package Manager (RPM) files. You must be running SLES 11 SP3 and a supported update version in order to update to the `Cray 7.2.UPxx` release.

Install the SMW Release Package

Follow these procedures to perform an initial, or clean installation of the Cray System Management Workstation (SMW) 7.2.UP03 release package.

NOTE: Only the procedure to install the SMW base operating system is different for rack-mount SMWs and desk-side SMWs. All other sections and procedures in this chapter apply to both types of SMW.

After completing the installation, if you need to configure, format, or zone the boot RAID system, see [Configure Boot RAID](#) on page 53.

Before You Begin

- Read the *SMW Release Errata* and the *SMW README* provided with your SMW release package for any additional installation-related requirements, corrections to this installation guide, and other relevant information about the release package.
- Read the Field Notices (FNs) related to kernel security fixes to identify any changes to this release package. Apply any needed changes before you install the software.
- For a system configured for SMW high availability (HA) with the SMW failover feature, there are installation differences for the two systems. Verify that you are using the correct installation guide: *Cray SMW Installation Guide (S-2480)* for systems with a single SMW; *SMW HA XC Installation Guide (S-0044)* for systems with high-availability (failover) SMWs.
- Verify that the network connections are in place (see [Network Connections](#) on page 6).
- Know which configuration values are site-specific and which are defaults (see [Configuration Values](#) on page 7).

Network Connections

The following network connections are required.

- A standalone SMW with a single quad-ethernet card has these private network connections:
 - eth0 - To the customer network
 - eth1 - To the Hardware Supervisory System (HSS) network
 - eth2 - Reserved for SMW failover
 - eth3 - To the boot node
- An SMW configured for SMW failover has a second quad-ethernet card with these connections:
 - eth4 - Used for SMW failover
 - eth 5 - Used for mirrored storage for the power management database (PMDB), if available

eth 6-7 - Reserved for future use

NOTE: Ethernet port assignments are valid only after the SMW software installation completes (see [Rack-mount SMW: Install the base operating system](#) on page 9).

- The SMW must have a Fibre Channel connection to the boot RAID.
- The boot node must have a Fibre Channel connection to the boot RAID.
- The service database (SDB) nodes must have Fibre Channel connections to the boot RAID.

IMPORTANT: Ensure you have protective covers for the Fibre Channel fiber optic cables and connectors. The SMW must be disconnected from the boot RAID before you perform the initial installation of the SLES software.

Configuration Values

The following IP addresses are set by default and are not site dependent.

Table 1. Default IP Addresses

IP Address	Description
10.1.0.1	Primary boot RAID controller
10.1.0.2	Secondary boot RAID controller
10.1.0.15	Storage RAID controller
10.1.1.1	SMW, eth1
10.2.1.1	SMW, eth2 - Reserved for SMW failover
10.3.1.1	SMW, eth3
10.3.1.254	boot node
10.4.1.1	SMW, eth4 - Reserved for SMW failover
127.0.0.1	localhost (loopback)

The following configuration values are site dependent. You may find it helpful to record the actual values for your site.

Table 2. Site-dependent Configuration Values

Description	Example	Actual Value
SMW hostname	xtsmw	
Domain	cray.com	
Aliases	cray-smw smw01	
Customer network IP address	192.168.78.68	
For rack-mount SMW only: iDRAC hostname	cray-drac	

Description	Example	Actual Value
For rack-mount SMW only: iDRAC IP address	192.168.78.69	
For rack-mount SMW only: iDRAC Subnet Mask	255.255.255.0	
For rack-mount SMW only: iDRAC Default GW	192.168.78.1	
Customer network netmask	255.255.255.0	
Default gateway	192.168.78.1	
Domain names to search	us.cray.com mw.cray.com	
Nameserver IP address	10.0.73.30 10.0.17.16	
Timezone	US/Central	
NTP servers	ntpghost1 ntpghost2	
X dimension	1-64	
Y dimension	Cray XE and Cray XK systems: 1-16; Cray XC Series Systems: 1-32	
Topology Class	Cray XE and Cray XK systems: 0, 1, 2, 3; Cray XC Series Systems: 0, 2 NOTE: Regardless of the number of cabinets in the system, Cray XC Series air cooled systems must be set to 0. Cray XC Series liquid cooled systems can be class 0 or 2.	

Passwords

The following default account names and passwords are used throughout the SMW software installation process. Cray recommends that you change these default passwords after completing the installation.

Table 3. Default System Passwords

Account Name	Password
root	initial0
crayadm	crayadm
cray-vnc	cray-vnc
mysql	None; you must create a password
admin (DDN™ boot RAID)	password

Account Name	Password
user (DDN boot RAID)	password
admin (DDN storage RAID)	password
user (DDN storage RAID)	password
root (iDRAC)	initial0

Back Up Your Current Software

Before you install the new software, back up your current SMW installation:

- If you have a rack-mount SMW, see [Rack-mount SMW: Create an SMW bootable backup drive](#) on page 33.
- If you have a desk-side SMW, see [Desk-side SMW: Create an SMW Bootable Backup Drive](#) on page 43.

Rack-mount SMW: Install the base operating system

1. If the SMW is up, `su` to `root` and shut it down.

```
crayadm@smw:~> su - root
smw:~ # shutdown -h now;exit
```

2. Disconnect the SMW connection to the boot RAID; disconnect the data cables and place protective covers on the fiber optic cables and connectors (if present).
3. Eject all the disk drives except for the primary disk in slot 0.
4. Power up the SMW. When the BIOS power-on self-test (POST) process begins, quickly press the F2 key after the following messages appear in the upper-right of the screen.

```

      F2 = System Setup
      F10 = System Services
      F11 = BIOS Boot Manager
      F12 = PXE Boot

```

When the **F2** keypress is recognized, the `F2 = System Setup` line changes to `Entering System Setup`.

After the POST process completes and all disk and network controllers have been initialized, the BIOS set-up screen appears.

5. Use the down-arrow key to highlight **Boot Settings**. Press the **Enter** key.

A window listing the following appears:

```

Boot Mode ..... BIOS
Boot Sequence ..... <ENTER>
USB Flash Drive Emulation Type..... <ENTER>
Boot Sequence Retry ..... <Disabled>

```

6. Use the down-arrow key to highlight **Boot Sequence**. Press the **Enter** key.

A window listing the following appears.

```

√ 1. Hard Drive C: (Integrated SAS 500 ID0A LUN0 ATA)
√ 2. Virtual CD
√ 3. Sata Optical Drive
√ 4. Embedded NIC 1 MBA v6.0.11 Slot 0100
√ 5. Virtual Floppy

```

7. Using the up-arrow and down-arrow keys to select, and using the space key to enable/disable entries, modify the list so that only the 1. Hard Drive C: entry has a check mark.

```

√ 1. Hard Drive C: (Integrated SAS 500 ID0A LUN0 ATA)
  2. Virtual CD
  3. Sata Optical Drive
  4. Embedded NIC 1 MBA v6.0.11 Slot 0100
  5. Virtual Floppy

```

8. Press the **Esc** key to exit the Boot Sequence window.
9. Press the **Esc** key again to exit the **Boot Settings** window.
10. Insert the base operating system DVD labeled `Cray-SMWbase11SP3-yyyymmdd` into the CD/DVD drive. (The DVD drive on the front of the SMW may be hidden by a removable decorative bezel.)
11. Press the **Esc** key a final time to save changes and exit the BIOS set up. A screen listing exit options appears.

```

Save changes and exit
Discard changes and exit
Return to Setup

```

12. Ensure that `Save changes and exit` is highlighted. Then press the **Enter** key. The SMW resets automatically.
13. When the BIOS POST process begins again, within 5 seconds, press the **F11** key after the following messages appear in the upper-right of the screen.

```

      F2 = System Setup
      F10 = System Services
      F11 = BIOS Boot Manager
      F12 = PXE Boot

```

When the **F11** keypress is recognized, the `F11 = BIOS Boot Manager` line changes to `Entering BIOS Boot Manager`.

14. Watch the screen carefully as text scrolls until you see `iDRAC6 Configuration Utility 1.57`. Within 5 seconds, press the **Ctrl-E** key when the prompt `Press <Ctrl-E> for Remote Access Setup within 5 sec...` displays.

```

  0   5   0 ATA      WDC WD5000BPVT-0 1A01      465 GB
LSI Corporation MPT2 boot ROM successfully installed!

```

```
iDRAC6 Configuration Utility 1.57
Copyright 2010 Dell Inc. All Rights Reserved
```

```
iDRAC6 Firmware Revision version: 1.54.15
Primary Backplane Firmware Revision 1.07
```

```
-----
                IPv6 Settings
-----
```

```
IPv6 Stack      : Disabled
Address 1       : ::
Default Gateway : ::
```

```
-----
                IPv4 Settings
-----
```

```
IPv4 Stack      : Enabled
IP Address      : 172. 31. 73.142
Subnet mask     : 255.255.255. 0
Default Gateway : 172. 31. 73. 1
Press <Ctrl-E> for Remote Access Setup within 5 sec...
```

The **iDRAC6 Configuration Utility** window appears.

15. In the **iDRAC6 Configuration Utility**, select **Virtual Media Configuration**.

The **Virtual Media Configuration** window appears.

16. Select the **Virtual Media** line until it indicates **Detached**.

```
Virtual Media ..... Detached
```

17. Press the **Esc** key to exit the **Virtual Media Configuration** window.

18. Press the **Esc** key again to exit the **iDRAC6 Configuration Utility** window.

The **BIOS Boot Manager** screen appears.

19. Use the up-arrow and down-arrow keys to highlight the **SATA Optical Drive** entry.

20. Press the **Enter** key to boot from the installation DVD.

21. Within 10 to 15 seconds after the SUSE Linux Enterprise Server boot menu displays, use the down-arrow key to scroll down and select the **Cray SMW Initial Install Rackmount** option.

```
- Boot from Hard Disk
- Cray SMW Initial Install
- Cray SMW Initial Install Rackmount
- Cray SMW Upgrade Install
- Repair Installed System
- Rescue System
- Check Installation Media
```

Then press the **Enter** key.

If you do not change the selection in time, the system will boot from the default selection, the hard disk. If that happens, shut down the SMW, then begin the power-up sequence again.

As the base installation progresses, the following phrases appear on your screen:

```
Analyzing Computer
System Probing
Preparing System for Automated Installation
Installation Settings
```

After these screens are displayed, the installation pauses on `Installation Settings`.

22. Partition the drive for installation of the base operating system and SMW software.

- a. Map the physical drive slots to the Linux device names.

IMPORTANT: Due to the nature of the hardware controllers on the SMW, the physical slots do not have a fixed logical mapping to the Linux `/dev/sd?` device names. Refer to the following table when identifying Linux `/dev` names to physical drive slots.

Table 4. Mapping of Linux `/dev` Names to Physical Drive Slots

Drive Slot	Linux <code>/dev</code> Name
0	<code>/dev/disk/by-path/pci-0000:05:00.0-sas-phy7-...</code>
1	<code>/dev/disk/by-path/pci-0000:05:00.0-sas-phy6-...</code>
2	<code>/dev/disk/by-path/pci-0000:05:00.0-sas-phy5-...</code>
3	<code>/dev/disk/by-path/pci-0000:05:00.0-sas-phy4-...</code>
4	<code>/dev/disk/by-path/pci-0000:05:00.0-sas-phy3-...</code>
5	<code>/dev/disk/by-path/pci-0000:05:00.0-sas-phy2-...</code>

Table 5. Physical Placement of Rack-mount SMW Chassis Hard Drives

Slot 0 (phy7)	Slot 2 (phy5)	Slot 4 (phy3)
Slot 1 (phy6)	Slot 3 (phy4)	Slot 5 (phy2)

The drives are assigned as follows:

- Slot 0: Target drive for installation of the base operating system and SMW software.
 - Slot 1: Will contain the backup installation of the base operating system and SMW software.
 - Slot 2: Will contain the database in the `/var/lib/mysql` directory.
 - Slot 3: Will contain the log, debug, and dump directories: `/var/opt/cray/log`, `/var/opt/cray/debug`, and `/var/opt/cray/dump`.
 - Slot 4: Optional drive for PMDB.
 - Slot 5:
- b. Remove and recreate the `phy7` partitions.
 1. For each device that is `phy7` Device Path (in this example, `/dev/sda1`) double-click on the Device name of that specific **Hard Disk** table entry, which opens the **Partition: `/dev/sd...`** table for that drive and which has **Edit**, **Resize** and **Delete** buttons below the table.

2. Select the **Delete** button. A pop-up window indicating **Really delete /dev/sd...** appears.
 3. Verify the device, and select **Yes** to delete the existing disk partition on that drive. The specific `/dev/sd...` entry will be removed from the displayed list.
 4. Repeat these steps for all entries associated with the `phy7` device path until the displayed list is empty.
- c. Create new partitions for `swap` and `root` for the `phy7` device sized for your system. This is the target boot device for the operating system installation.

1. Select the **Add** button at the bottom of the **Hard Disk: /dev/sd...** window, which opens the **Add Partition on /dev/sd...** screen.
2. Select **Primary Partition**, then select **Next**. The **New Partition Size** screen displays.
3. Select **Custom** size, then enter the swap partition size into the box. This swap partition size should be the same as the total system RAM memory on the SMW.

Enter either **8GB** or **32GB**, depending on the total system RAM memory on your SMW.

TIP: If you are unsure of the value to use: For an SMW that has one power supply, enter **8GB**. For an SMW that has two power supplies, enter **32GB**.

Then select **Next**. The **Formatting Options and Mounting Options** screen displays.

4. Select **Format partition** and select **Swap** from the pull-down menu under **File system**.
5. Select **Mount partition** and select **Swap** from the pull-down menu under **Mount Point**.
6. Select the **Finish** button. The **Hard Disk: /dev/sd...** window displays again, but with the new partition entry.
7. Select the **Add** button at the bottom of the window, which opens the **Add Partition on /dev/sd...** screen.
8. Select **Primary Partition**, then select **Next**. The **New Partition Size** screen displays.
9. Select **Maximum Size**, then select **Next**. The **Formatting Options and Mounting Options** screen displays.
10. Select **Format partition** and select **Ext3** from the pull-down menu under **File system**.
11. Select **Mount partition** and select **/** from the pull-down menu under **Mount Point**.
12. Select **Finish**. Partitioning of the target boot device for the operating system installation has been completed. On the **Expert Partitioner** screen, select **Accept** to accept the changes. A pop-window stating the following will be displayed:

```
Changes in disk partitioning were detected since the time the bootloader
was configured.
Do you want to proposed bootloader configuration again?
If yes, all previous bootloader configuration will be lost.
If not, you probably need to change the configuration manually.
```

Select **OK**. This pop-up window has a count-down timer that selects **OK** automatically if not interrupted with the **Stop** button. The display returns to the **Installation Settings** screen.

- d. Confirm the disk partitions.

A single `root` device and a single `swap` device, both associated with `phy7` are the only partitions listed. For example:

```
Partitioning
Create root partition /dev/sdc2 (146.92 GB) with ext3
use /dev/sdc1 as swap
```

To make additional, site-specific changes, select the **Partitioning** section header and return to the **Expert Partitioner** screen.

23. Confirm the language for the SMW. English (US) is the default language. To change the primary language select the **Language** heading in the **Installation Settings** screen. The **Languages** window opens. Select your language choice from the drop-down menu. You can also select multiple secondary languages. Then select **Accept** at the bottom of the window.
24. On the **Installation Settings** screen, select **Install** at the bottom of the screen. The **Confirm Installation** window appears. If you want to check or change settings, select **Back**; otherwise, select **Install** to confirm and to install the operating system.
The installation runs for approximately 30 minutes. The process automatically reboots the SMW from the hard disk, and the installation continues with system configuration.
25. The **Network Settings** window for configuring the customer network appears. Select the entry labeled **eth0 Customer Network Ethernet**. On the **OK** tab, select edit and do the following:
 - a. Enter the IP address.
 - b. Enter the subnet mask.
 - c. Enter the short hostname and select **Next**.
 - d. Select the **Hostname/DNS** tab. In the **Hostname/DNS and Name Server Configuration** window, enter the hostname, the domain name, the name server values, and the domain names to search. You must enter the hostname and domain name separately.

Host Name	Domain Name
smwhost	my.domain.com

If you enter a fully-qualified hostname that includes the domain name, the hostname is accepted but without the periods; for example, a hostname of `smwhost.my.domain.com` is converted to `smwhostmydomaincom`.

- e. Select the **Routing** tab. In the **Routing** window, enter the default gateway IP address. Then select **OK**.
26. The **Clock and Time Zone** window appears.
 - a. Select your time zone.
 - b. If necessary, adjust the time of day.
 - c. Verify that **Hardware Clock Set To UTC** is selected.
 - d. Select **OK**. At this point, the system finishes booting and enters multiuser mode.
The SMW base operating system, `Cray-SMWbase11SP3`, is now installed. Continue with the next steps.
27. Remove the protective covers from the fiber optic cables and connectors, clean the ends of the cables and connectors, and reconnect the data cables.

28. Reboot the SMW to allow the SMW to discover the drives properly.

```
smw:~ # reboot
```

29. Enable Integrated Dell™ Remote Access Controller 6 (iDRAC6). See [Enable an Integrated Dell Remote Access Controller \(iDRAC6\) on a rack-mount SMW](#) on page 103

Installation of the base operating system for the Rack-mount SMW is now complete. To install and configure the SMW software packages go to [Install and Configure the SMW Software Packages](#) on page 18.

Desk-side SMW: Install the base operating system

1. If the SMW is running, `su` to `root` and shut it down.

```
crayadm@smw:~> su - root
smw:~ # shutdown -h now;exit
```

2. Disconnect the SMW connection to the boot RAID; disconnect the data cables and place protective covers on the fiber optic cables and connectors (if present).
3. Power up the SMW. When the BIOS power-on self-test (POST) process begins, quickly press the **Delete** key after the following messages appear on the screen.

```
Press DEL to run Setup
Press F12 if you want to boot from the network
Press F11 for BBS POPUP
Initializing USB Controllers .. Done.
...
USB Device(s): 1 Keyboard, 1 Mouse
```

The last line in this example indicates that the SMW is ready to accept keyboard input.

After the POST process completes, the BIOS Setup Utility screen appears.

4. Use the right-arrow key to select the Boot menu. Use the down-arrow key to highlight **Boot Device Priority**. Press the **Enter** key.
5. Verify the boot priority list.

```
1st Boot Device [SATA:3M-ST3320613A]
2nd Boot Device [CD/DVD:PM-ATAPI DV]
3rd Boot Device [1st Floppy Drive]
```

If necessary, change the boot priority list so that the local hard drive is the first boot device and the local CD/DVD drive is the second boot device.

6. Press the **Esc** key to return to the BIOS Setup Utility screen.
7. Press the down-arrow key to select **Hard Disk Drives**, and press the Enter key.

8. Verify that the drive order shown corresponds to the physical placement of the hard drives in the SMW. The drives should be placed in the following order:

```
1st Drive    [SATA: 3M-ST3320613A]
2nd Drive    [SATA: 4M-ST3320620A]
3rd Drive    [SATA: 5M-ST3320613A]
```



CAUTION: If the drive order is set incorrectly, the installation process appears to continue normally, but it aborts when the kernel fails to load at the end of the SMW installation process. If this happens, the only solution is to begin the entire installation process anew.

9. Press the **Esc** key to return to the Boot menu. Then press the left-arrow key to select the Main menu. On the Main menu, the System Time option should already be selected. Set the system date and time to Coordinated Universal Time (UTC). Execute the command:

```
smw~> date -u
```

to determine the current UTC.

10. Insert the base operating system DVD labeled `Cray-SMWbase11SP3-yyyymmdd` into the CD/DVD drive.

11. Press the **F10** key to save the BIOS changes. A dialog box displays the question **Save configuration changes and exit setup?**

- Press **Enter** to save your changes and exit the setup utility. This will automatically reboot the SMW.
- Select **Cancel** to cancel acceptance of BIOS changes, in order to adjust a setting.

12. When the POST process begins again, quickly press the **F11** key to enter **Boot Device Selection** mode.

```
Press DEL to run Setup
Press F12 if you want to boot from the network
Press F11 for BBS POPUP
Initializing USB Controllers .. Done.
...
USB Device(s): 1 Keyboard, 1 Mouse
```

(The last line in this example indicates that the SMW is ready to accept keyboard input.)

13. At the **Please select boot device** prompt, use the up-arrow and down-arrow keys to select the CD/DVD drive, and press the **Enter** key. A blue-background Cray logo screen with a DVD-boot option menu appears.
14. Within 10 to 15 seconds after the SUSE Linux Enterprise Server boot menu displays, use the down-arrow key to scroll down and select the **Cray SMW Initial Install** option, then press **Enter**.

```
- Boot from Hard Disk
- Cray SMW Initial Install
- Cray SMW Initial Install Rackmount
- Cray SMW Upgrade Install
- Repair Installed System
- Rescue System
- Check Installation Media
```

As the base installation progresses, the following status messages appear on your screen:

```
Analyzing Computer
System Probing
```

Preparing System for Automated Installation Installation Settings

After these screens are displayed, the installation pauses on **Installation Settings**.

15. Confirm which disk partitions are marked for use as swap space. Use only the disk partition on the primary SMW disk as swap for the SMW. Check the **Partitioning** heading. If any disk has a swap partition on it and there is a line that says that device is used as swap, it must be changed. For example:

```
- Use /dev/sdb1 as swap
- Use /dev/sdd3 as swap
```

A line similar to the following shows the primary SMW disk partition to be used for swap:

```
- Create swap partition /dev/sda1
```

If any other partitions are marked as `Use /dev/xxxx as swap`, you must continue with the steps below to change their designation so that they are not used as swap.

- a. To designate that a disk partition not be used for swap on the SMW, select the **Partitioning** heading, which opens the **Expert Partitioner** window.
 - b. Select the **+** sign next to **Hard Disks** to display the hard disks that are connected.
 - c. Select the **+** sign next to a disk device that incorrectly has a `swap` partition in the **Mount Point** column.
 - d. Select **sdb1 (disk device)** to open the **Edit Partition** window
 - e. Select **Edit** at the bottom of the window.
 - f. Select **Do not mount partition**.
 - g. Select **Finish** in the **Edit Partition** window.
 - h. Select **Accept**.
16. Confirm the language for the SMW. English (US) is the default language. To change the primary language, select the **Language** heading in the **Installation Settings** window. The **Languages** window opens. Select your language choice from the drop-down menu. You can also select multiple secondary languages. Then select **Accept** at the bottom of the window to continue the installation.
 17. On the **Installation Settings** window, select **Install** at the bottom of the window. The **Confirm Installation** window appears. If you want to check or change settings, select **Back**; otherwise, select **Install** to confirm and to install the operating system. The installation script takes approximately 30 minutes to complete. The process automatically reboots the SMW from the hard disk, and the installation continues with system configuration.
 18. The **Network Settings** window for configuring the network appears.
 - a. Select the entry labeled **eth0 Customer Network Ethernet**.
 - b. On the **Overview** tab, select **Edit** and enter the IP address, the subnet mask, and the short hostname. Select **Next**.
 - c. Select the **Hostname/DNS** tab. In the **Hostname/DNS and Name Server Configuration** window, enter the hostname, the domain name, the name server values, and the domain names to search. You must enter the hostname and domain name separately. If you enter a fully-qualified hostname that includes the

domain name, the hostname is accepted but without the periods; for example, a hostname of `smwhost.my.domain.com` is converted to the incorrect `smwhostmydomaincom`

Host Name	Domain Name
smwhost	my.domain.com

- d. Select the **Routing** tab. In the **Routing** window, enter the Default Gateway IP address. Then select **OK**.

19. The **Clock and Time Zone window appears.**

- a. Select your time zone.
- b. Adjust the time of day, if necessary.
- c. Ensure that **Hardware Clock Set To UTC** is selected.
- d. Select **OK**

At this point, the system comes up the rest of the way to multiuser mode for login on the console. The SMW base operating system is now installed. Reconnect the system:

- Remove the protective covers from the fiber optic cables and connectors, clean the ends of the cables and connectors, and reconnect the data cables.
- Seat the previously ejected disk drives.
- Reboot the SMW to allow the SMW to discover the drives properly.

```
smw:~ # reboot
```

Install and Configure the SMW Software Packages



CAUTION: You must shut down your Cray system before you install the Cray SMW software packages.

Shut down the Cray system

1. Log on to the SMW as `crayadm` and confirm the Cray system is shut down.

```
crayadm@smw:~> ping boot
```

If the command responds with "alive", then it is up and needs to be shut down.

2. Shut down the system by typing the following command.

```
crayadm@smw:~> xtbootsys -s last -a auto.xtshutdown
```

For more information about using automation files, see the `xtbootsys(8)` man page.

Install and configure SMW software

This procedure takes approximately 30 minutes.

1. Log on to the SMW as `crayadm`, open a terminal window, and `su` to `root`.

```
crayadm@smw:~> su - root
smw:~ #
```

2. If the base operating system DVD (`Cray-SMWbase11SP3-yyyymmdd`) is still in the CD/DVD drive, eject it.

```
smw:~ # eject
```

3. Place the `Cray 7.2UPxx Software` DVD in the drive and mount it.

```
smw:~ # mkdir -p /media/cdrom
smw:~ # mount /dev/cdrom /media/cdrom
```

NOTE: If you experience problems mounting the DVD after the initial Linux installation, reboot the SMW and perform this procedure again, from the beginning.

4. Copy the `SMWinstall.conf` file from `/media/cdrom` to `/home/crayadm`.

```
smw:~ # cp /media/cdrom/SMWinstall.conf /home/crayadm
```

5. Change the permissions to make the file writable by the user only.

```
smw:~ # chmod 644 /home/crayadm/SMWinstall.conf
```

6. Edit the `SMWinstall.conf` file to customize it for your site.

```
smw:~ # vi /home/crayadm/SMWinstall.conf
```

The `SMWinstall.conf` file contains settings for:

- the system interconnection network type, e.g., Aries or Gemini
 - the name of the local Network Time Protocol (NTP) servers
 - configuring SMWs that have additional disks
 - configuring an SMW for a Cray XE6m, Cray XK6m, or Cray XK7m mid-range system
 - configuring the Cray Lightweight Log Manager (LLM)
 - enabling or disabling the Linux `sar` service on the SMW
- a. Specify the system interconnection network type for your system. There is no default setting.
 - b. By default, the Linux `sar` service on the SMW is enabled. You may choose to disable it.
 - c. By default, controller log forwarding is enabled. You may choose to disable it.
 - d. For rack-mount SMWs: Set the `LOGDISK`, `DBDISK`, and `PMDISK` variables. Note that the disk device specified by the `LOGDISK` variable must have more disk space than the device specified by the `DBDISK` variable.
 1. Use the `LOGDISK` variable to specify the disk device name to be used for logging (`/var/opt/cray/log`, `/var/opt/cray/dump`, and `/var/opt/cray/debug`). The entire disk will be formatted for this use. On a rack-mount SMW, this will be the `phy4` disk.

IMPORTANT: For an SMW HA system, define this variable as if the system was a stand-alone SMW. The shared storage for logging will be configured later in the SMW HA configuration process.

2. Use the `DBDISK` variable to specify the disk device name to be used for the SMW HSS database (`/var/lib/mysql`). The entire disk will be formatted for this use. On a rack-mount SMW, this will be the `phy5` disk.

IMPORTANT: For an SMW HA system, define this variable as if the system was a stand-alone SMW. The shared storage for the SMW HSS database will be configured later in the SMW HA configuration process.

3. (Optional) Use the `PMDISK` variable to specify the disk device name to be used for the PostgreSQL Database (`/var/lib/pgsql`). The entire disk will be formatted for this use. On a rack-mount SMW this will be the `phy3` disk.

IMPORTANT: For an SMW HA system that will use shared or mirrored storage for `/var/lib/pgsql`, the `PMDISK` variable is not used. The disk device for `/var/lib/pgsql` is set up later in the SMW HA configuration process. Do not set the `PMDISK` or `PMMOUNT` variables.

The `LOGDISK`, `DBDISK`, and `PMDISK` variables must be persistent disk device names. To use the default persistent disk device names, remove the comment character (`#`) from the following lines in the file:

```
#LOGDISK=/dev/disk/by-path/pci-0000:05:00.0-sas-phy4-0x4433221104000000-lun-0
#LOGMOUNT=/var/opt/cray/disk/1

#PMDISK=/dev/disk/by-path/pci-0000:05:00.0-sas-phy3-0x4433221103000000-lun-0
#PMMOUNT=/var/lib/pgsql

#DBDISK=/dev/disk/by-path/pci-0000:05:00.0-sas-phy5-0x4433221105000000-lun-0
#DBMOUNT=/var/lib/mysql
```

If you have a nonstandard configuration, modify the persistent disk device names accordingly.

- e. For Cray XE6m, Cray XK6m, or Cray XK7m systems: Set the `MCLASS` and `NumChassis` variables. These configuration settings do not apply to Cray XE6m-200, Cray XK6m-200, or Cray XK7m-200 systems.
 1. For a Cray XE6m, Cray XK6m, or Cray XK7m system, set the `MCLASS` variable to `TRUE`. For all other Cray systems, including Cray XE6m-200, Cray XK6m-200, and Cray XK7m-200 systems, `MCLASS` must retain the default setting of `FALSE`.
 2. For a Cray XE6m, Cray XK6m, or Cray XK7m system only, remove the comment character from the `NumChassis` line. The `NumChassis` variable indicates the number of chassis that are in the Cray `MCLASS` system. Choices are 1, 2, 3, 4, 6, 8, 9, 12, 15, or 18.

7. Execute the `SMWinstall` installation script, which updates the base operating system software with SMW security updates and SMW software.

```
smw:~ # /media/cdrom/SMWinstall
```

The output of the installation script is displayed on the console. The `SMWinstall` installation script also creates log files in `/var/adm/cray/logs`.

If for any reason this script fails, it can be rerun without adverse side effects. However, rerunning this script can generate numerous error messages as the script attempts to install already-installed RPMs. You can safely ignore these particular messages.

8. By default, the SMWinstall program enables LLM. During the SMW software installation or upgrade the SMWinstall program automatically updates the `llm.conf` configuration file. If you made changes to the `llm.conf` file provided with the release package, you must reapply your changes after an SMW software installation or upgrade. As part of the update process, the `llm.conf` configuration file is saved as `/etc/opt/cray/llm/llm.conf.rpmsave`.

9. Reboot the SMW.

```
smw:~ # reboot
```

10. To set up a Virtual Network Computing (VNC) server for remote access to the SMW, see [Enable Remote Access to the SMW](#) on page 97.

From this point on, whenever the SMW presents a session selection menu, always select SUSE Linux Enterprise Server 11 SP3.



CAUTION: Do not execute any SMW commands until after you have rebooted your SMW and completed the procedures in this section.

Discover the Cray System Hardware and Power Up the Full System

IMPORTANT: If an installation step fails because of a hardware issue, such as a cabinet failing to power-up, when that issue is resolved you must go back to the last successful step in the installation procedure and continue from there. Do not skip steps or continue out of order.

Bootstrap Hardware Discovery

1. Log on to the SMW as `crayadm` and `su` to `root`.

```
crayadm@smw:~> su - root
smw:~ #
```

2. **RESTRICTION:** For a Cray XC Series system Aries™ network SMW only; if installing software on a Cray XE or a Cray XK system (Gemini™ network) skip this step.

Update the controller boot image.

The version used in the command argument for `hss_make_default_initrd` should match the version specified in the `lsb-cray-hss` line in output from the `crms-release` file. This directory will not exist until the `hss_make_default_initrd` command completes.

```
smw:~ # cat /opt/cray/hss/default/etc/crms-release
HSS-CRMS Wed May 15 12:13:32 CDT 2013 on hssbld0 by bwdev
lsb-cray-hss-7.2.0-1.0702.31509.447
smw:~ # hss_make_default_initrd /opt/cray/hss-images/master/7.2.0-1.0702.31509.447
::: Verifying base RPM list to the manifest
(additional status messages)
::: Removing unwanted files from the root
```

```

=====
The new initrd hierarchy is now in /opt/cray/hss-images/
master/ 7.2.0-1.0702.32944.237.
Running hssclone.
Image Clone Complete: /opt/cray/hss-images/default
Running hsspackage.
. . .(additional status messages). . .
inking /opt/cray/hss-images/default/HSS32/initrd.img /opt/tftpboot/initrd.img

```

The `xtdiscover` command needs to collect some basic information in order to bootstrap the hardware discovery process. In the next step, you should be prepared to input the following information:

```

Abort or continue [a or c]
Network type (g=Gemini, a=Aries, q=quit)
Single-Slot Tester [y or n]
Maximum X cabinet size [1-64]
Maximum Y cabinet size [1-16]
Network topology class [0-3] (0 or 2 for Cray XC30 systems, 0 for Cray XC30-AC
systems)
Boot node name [c0-0c0s0n1]
SDB node name [c0-0c0s2n1] (Cray XE or Cray XK systems)
SDB node name [c0-0c0s1n1] (Cray XC30 systems)
Pathname to the default boot image [/raw0]

```

- The default boot node provided by `xtdiscover` is `c0-0c0s0n1` for Cray XE or Cray XK systems (Gemini™ network) and Cray XC Series systems (Aries™ network).
- The default SDB node provided by `xtdiscover` is `c0-0c0s2n1` for Cray XE or Cray XK systems (Gemini™ network) and Cray XC Series systems (Aries™ network).
- By default, the `xtdiscover` command uses the `/opt/cray/hss/default/etc/xtdiscover.ini` file. When you perform an SMW update, the process will not overwrite the `xtdiscover.ini` file; the new version of the file will be placed in `xtdiscover.ini.dist`.

3. Run the `xtdiscover --bootstrap` command:

```

smw:~ # xtdiscover --bootstrap
***** xtdiscover started *****
Command line: xtdiscover --bootstrap
USER: crayadm pts/0      Jun 14 13:14 (mycomputer.company.com)
Using ini file '/opt/cray/hss/default/etc/xtdiscover.ini'
xtdiscover is about to discover new hardware.
This operation may significantly modify the system database.
Please enter 'c' to continue, or 'a' or 'q' to abort [c]: c
Please enter network type (g=Gemini, a=Aries, q=quit): g
Is this system a Single-Slot Tester? y/n, q=quit [n]: n
Setting system type to Gemini
Discovering Gemini-based system...
Enter maximum X cabinet size (columns) [1-64], q=quit: 1
Enter maximum Y cabinet size (rows) [1-16], q=quit: 1
Adding hosts and routes for 1 cabinet...done.
Enter your system's network topology class [0]: 0 (0 for Cray XC30 systems, 2
for Cray XC30-AC systems)
Setting topology class to 0
Suspending State Manager for discovery phase 1...
Suspend successful.
Saving current configuration...done.
Force new config selected.  Initializing datastore...done.
Discovering cabinets:
[1 out of 1]

```

```

Found 1 cabinet.
xtdiscover will create a single system partition (p0)
containing all discovered cabinets. If you need to create
additional partitions, use 'xtcli part_cfg add'.
Enter the boot node name [c0-0c0s0n1]: c0-0c0s0n1
Enter the SDB node name [c0-0c0s2n1]: c0-0c0s2n1
Enter the absolute pathname to the default boot image [/raw0]: /raw0
Storing base cabinet data...done.
Resuming State Manager for L1 wipe and reboot...Resume successful.
. . .[additional status messages]. . .
Discovery Phase 1 of 3 complete.

```

NOTE: Prior to powering on the cabinets, xtdiscover prompts you to disable any blades that should not be powered on. After you have disabled these blades, confirm that you are ready to continue or abort (the xtdiscover prompt is Continue or abort [a or c]). The xtdiscover command then proceeds without any further questions.

```

xtdiscover is about to power on the cabinets.
*** IF YOU NEED TO DISABLE COMPONENTS TO AVOID THEM
*** BEING POWERED ON, PLEASE DO SO NOW USING 'xtcli disable'

Please enter 'c' to continue, or 'a' or 'q' to abort [c]: c
Suspending State Manager for discovery phase 2...
Suspend successful.
. . .[additional status messages]. . .
Done.
Discovery complete
***** xtdiscover finished *****

```

NOTE: If the xtdiscover command fails with the message, The following cabinets were not detected by heartbeat:, power cycle the cabinet controller, and retry the xtdiscover -- bootstrap command.

RESTRICTION: The next three steps are for a Cray XC Series system (Aries™ network) SMW only; if installing software on a Cray XE or a Cray XK system (Gemini™ network) skip to Step 7.

4. Power down the system.

```
smw:~ # xtcli power down s0
```

5. Reboot the cabinet controllers, then ensure that all cabinet controllers are up.

```

smw:~ # xtccreboot -c all
xtccreboot: reboot sent to specified CCs
smw:~ # xtalive -l cc

```

6. Power up the system.

```
smw:~ # xtcli power up s0
```

RESTRICTION: The next three steps are for a Cray XE or a Cray XK system (Gemini network) SMW only. For Cray XC Systems, the bootstrap process is now complete.

7. Exit root.

```
smw:~ # exit
```

8. Log on as `crayadm`, and execute the `xtflash` command to update the L0s and L1s with the current firmware.

```
crayadm@smw:~> xtflash s0
It took 0 seconds for 'ping' to complete to all L1s.
It took 0 seconds for L1s to become available.
It took 0 seconds for 'ping' to complete to all L0s.
It took 0 seconds for L0s to become available.
spawn fm -qRM -t l1 s0
It took 2 seconds for 'query' to complete to all L1s.
xtrsh -m "-[0-7]$" -f /tmp/xtrsh.Sk5yg9Uz -l root -s "if [ -x /etc/loadnor ]; \
then /etc/loadnor; fi ; if md5sum /dev/mtd/4 | grep
fec74d6d797dd8b36e68282ad43d0773; \
then echo OLDBIOS; fi"
There is one L1 not up-to-date.
spawn fm -qRM -t l0 s0
It took 2 seconds for 'query' to complete to all L0s.
xtrsh -m s -f /tmp/xtrsh.6ef31uGH -l root -s "if [ -x /etc/loadnor ]; \
then /etc/loadnor; fi ; if md5sum /dev/mtd/4 | grep
fec74d6d797dd8b36e68282ad43d0773; \
then echo OLDBIOS; fi"
There are 24 L0s not up-to-date.
#####
# Attempt #1 #
#####
spawn fm -RM -t l1 c0-0
It took 157 seconds (2 minutes, 37 seconds) for 'flash' to complete to all L1s.
There is one L1 not up-to-date.
xtrsh -m "-[0-7]$" -f /tmp/xtrsh.b0zXZERW -l root -s "reboot"
Sleeping for 180 seconds...
It took 0 seconds for 'ping' to complete to all L1s.
It took 0 seconds for L1s to become available.
spawn fm -RM -t l0 c0-0c0s0 c0-0c0s1 c0-0c0s2 c0-0c0s3 c0-0c0s4 c0-0c0s5
c0-0c0s6 c0-0c0s7 \
c0-0c1s0 c0-0c1s1 c0-0c1s2 c0-0c1s3 c0-0c1s4 c0-0c1s5 c0-0c1s6 c0-0c1s7
c0-0c2s0 c0-0c2s1 \
c0-0c2s2 c0-0c2s3 c0-0c2s4 c0-0c2s5 c0-0c2s6 c0-0c2s7
It took 180 seconds (3 minutes, 0 seconds) for 'flash' to complete to all L0s.
There are 24 L0s not up-to-date.
.
.
.
Everyone is properly flashed.
It took 871 seconds (14 minutes, 31 seconds) to run this application.
crayadm@smw:~>
```

If all the L0s are running a compatible BIOS firmware, then `xtflash` completes successfully, and you can continue to [Discover the Cray System Hardware](#) on page 25. However, if the L0s do not get flashed with the updated BIOS firmware, `xtflash` will halt and display a list of those L0s that need to be updated. These L0s must have their BIOS reflashed before `xtflash` will continue. Perform the BIOS flash only on the components that require it.

TIP: If `xtflash` halts and displays a list of those L0s that need to be updated, Cray recommends reflashing the Cray L0s in groups of a few cabinets at a time, rather than the entire system at once. If there is a catastrophic failure during the BIOS flashing procedure, the only way to recover the BIOS is to reflash each BIOS chip by hand. By flashing only a few cabinets at a time, you limit the risk that a catastrophic failure can cause. In addition, if you flash the BIOS on an L1/L0 that already has a compatible BIOS (like an SIO), then you must power cycle the L1/L0 after flashing it.

In this example, both `c0-0c1s3` and `c0-0c1s4` need to have their BIOS chips reflashed. Reflash their BIOS chips by using the `fm` command with the BIOS option `-B` and a space-separated or comma-separated list of components:

```
crayadm@smw:~> fm -B -t 10 c0-0c1s3 c0-0c1s4
100% complete
c0-0c1s3: ok
c0-0c1s4: ok
```

The `fm` command also allows you to specify specific chassis or cabinets. For example, to reflash all the L0 BIOS chips on chassis `c0-0c1` and cabinets `c1-0` and `c2-0`, enter:

```
crayadm@smw:~> fm -B -t 10 c0-0c1 c1-0 c2-0
```

9. After the BIOS chips have been reflashed, execute the `xtflash` command again to continue on with reflashing the L1 and L0 Linux firmware.

```
crayadm@smw:~> xtflash s0
```

The bootstrap process is now complete. The next task is to discover the Cray system hardware.

Discover the Cray System Hardware

1. To detect the Cray system hardware components on the system, run the `xtdiscover` command. This command creates entries in the system database to describe the hardware. After running `xtdiscover`, you can use the `xtcli` command to display the configuration. For more detailed information, see the `xtdiscover(8)` man page.

The `xtdiscover` command collects some basic information in order to bootstrap the hardware discovery process, warns that changes will be made, and then asks for a confirmation of aborting or continuing. You will need to provide the X and Y cabinet sizes as well as topology from [Site-dependent Configuration Values](#) in [Configuration Values](#) on page 7 for `xtdiscover` to accurately discover the hardware. The `xtdiscover` command may prompt you to execute the `xtbounce` command in a separate window and then continue.

An example of the output of the `xtdiscover` and `xtbounce` commands follows:

```
crayadm@smw:~> su - root
smw:~ # xtdiscover
***** xtdiscover started *****
Command line: xtdiscover
USER: crayadm pts/0          Jun 14 13:15 (mycomputer.company.com)

Using ini file '/opt/cray/hss/default/etc/xtdiscover.ini'

xtdiscover is about to discover new hardware.
This operation may significantly modify the system database.

Please enter 'c' to continue, or 'a' or 'q' to abort [c]: c

Please enter network type (g=Gemini, a=Aries, q=quit): g

Is this system a Single-Slot Tester? y/n, q=quit [n]: n

Setting system type to Gemini
```

Discovering Gemini-based system...

Enter maximum X cabinet size (columns) [1-64], q=quit: 1
 Enter maximum Y cabinet size (rows) [1-16], q=quit: 1
 Adding hosts and routes for 1 cabinet...done.

Enter your system's network topology class [0]: 0 (0 for Cray XC30 systems, 2 for Cray XC30-AC systems)
 Setting topology class to 0

Suspending State Manager for discovery phase 1...
 Suspend successful.
 Saving current configuration...done.

Discovering cabinets:
 [1 out of 1]
 Found 1 cabinet.

xtdiscover will create a single system partition (p0) containing all discovered cabinets. If you need to create additional partitions, use 'xtcli part_cfg add'.

Enter the boot node name [c0-0c0s0n1]: c0-0c0s0n1
 Enter the SDB node name [c0-0c0s2n1]: c0-0c0s2n1
 Enter the absolute pathname to the default boot image [/raw0]: /raw0

Gathering base cabinet attributes:
 [1 out of 1]
 Finished gathering cabinet attributes.

Clearing database...done.
 Verifying phase 1 configuration...done.
 Storing base cabinet data...done.
 Resuming State Manager for power-up and bounce...Resume successful.

Discovery Phase 1 of 3 complete.

xtdiscover is about to power on the cabinets.
 *** IF YOU NEED TO DISABLE COMPONENTS TO AVOID THEM
 *** BEING POWERED ON, PLEASE DO SO NOW USING 'xtcli disable'

Please enter 'c' to continue, or 'a' or 'q' to abort [c]: c

Suspending State Manager for discovery phase 2...
 Suspend successful.

Loading base component data for discovery phase 2...done.

Powering on cabinet...
 1 cabinet will be powered on:
 [1 out of 1]
 Cabinets powered on.

Discovering component phase 2 (blade) state:
 [24 out of 24]
 Finished discovering component phase 2 (blade) state.

Discovering component phase 2 (blade) attributes:
 [24 out of 24]
 Finished discovering component phase 2 (blade) attributes.
 Verifying phase 2 configuration...done.

```
Summary of blades discovered:
Total:    24    Service:    11    Empty:    0    Disabled:    0
```

```
Storing attribute data...done.
```

```
Discovery Phase 2 of 3 complete.
```

```
Blades changed: 24; blades needing to be bounced: 24.
```

```
Resuming State Manager for bounce...Resume successful.
```

```
24 blades should be bounced using the command
in file /opt/cray/hss/default/etc/xtdiscover-bounce-cmd
```

```
In a separate window, please bounce the system now to continue discovery.
```

```
After bounce completes, enter 'c' to complete discovery
or 'q' or 'a' to abort [c]:
```

2. In a separate window, as `crayadm`, bounce the system using the `/opt/cray/hss/default/etc/xtdiscover-bounce-cmd` command.

```
crayadm@smw:~> /opt/cray/hss/default/etc/xtdiscover-bounce-cmd
```

3. After the `xtbounce` has finished, return to the primary `xtdiscover` window and respond to the following prompt:

```
In a separate window, please bounce the system now to continue discovery.
```

```
After bounce completes, enter 'c' to complete discovery
or 'q' or 'a' to abort [c]: c
```

```
Suspending State Manager for discovery phase 3...
Suspend successful.
```

```
Discovering component phase 3 (blade/node) attributes:
[24 out of 24]
```

```
Finished discovering component phase 3 attributes.
```

```
Verifying configuration...
```

```
INFO: 24 newly discovered components were added.
```

```
INFO: 24 components in previous configuration were deleted.
```

```
INFO: Added the following hardware:
```

```
24 slots
```

```
INFO: Removed the following hardware:
```

```
24 slots
```

```
INFO: Configuration change details are in
/opt/cray/hss/default/etc/xtdiscover-config-changes.diff
verification complete.
```

```
Discovery complete.
```

```
Commit to Database? [y]es, [n]o (repeat stage 3), [a]bort: y
```

```
Add discovered power management node descriptors.
```

```
Clearing database...done.
```

```
Storing component attribute data to database...done.
```

```
Storing component attribute data to fast file...done.
```

```
Updating component location history...done.
```

```
Restarting RSMS daemons for normal operation:
```

```
Stopping RSMS services: xtremoted erfsd xtpmd cm sedc_manager bm nm sm erdh
```

```
INFO:
```

```
connection to event router lost
```

```

INFO: unable to connect to event router, sleeping ...
erd
done
Starting RSMS services: erd erdh sm nm bm sedc_manager cm
Waiting for state manager
Starting RSMS services: xtpmd erfds xtremoted
Flushing and installing cabinet routes...RTNETLINK answers: File exists
done.
Partition p0 is active but not booted; skipping background commands.
done
Done.
Discovery complete
Waiting for SM ready notification.
Received SM ready notification.
***** xtdiscover finished *****

```

4. **RESTRICTION:** For a Cray XC30 system (Aries network) SMW only. Skip these substeps if installing on a Cray XE or XK system.

Exit from `root`.

```
smw:~ # exit
```

- a. Run the `rtr --discover` command to determine the exact configuration of the HSN.

```
crayadm@smw:~> rtr --discover
```

If the system was not bounced previously you may see the following message:

```
System was not bounced in diagnostic mode, should I re-bounce? Continue (y/n)?
```

If so, respond with `y`

- b. Update the firmware. Execute the `xtzap` command to update the components.



CAUTION: The `xtzap` command is normally intended for use by Cray Service personnel only. Improper use of this restricted command can cause serious damage to the computer system.

```
crayadm@smw:~> xtzap -r -v s0
```

IMPORTANT: The Cray XC30 system also requires an update to the NVIDIA® BIOS (nvBIOS) for the NVIDIA K20X graphics processing units (GPUs). This update is done after CLE has been booted. For more information, see *Installing and Configuring Cray Linux Environment (CLE) Software (S-2444)*.

- c. If the output of `xtzap` includes a "Revision Mismatches" section, some firmware is out of date and needs to be reflashed. To update, run `xtzap` with one or more of the options described in the next paragraph.

While the `xtzap -a` command can be used to update all components with a single command, it may be faster to use the `xtzap -blade` command when only blade types need to be updated, or the `xtzap -t` command when only a single type needs to be updated. On larger systems, this can be a significant time savings.

This is the list of all cabinet level components:

```
cc_mc (CC Microcontroller)
cc_bios (CC Tolapai BIOS)
```

```
cc_fpga (CC FPGA)
chia_fpga (CHIA FPGA)
```

This is a list of all blade level components:

```
cbb_mc (CBB BC Microcontroller)
ibb_mc (IBB BC Microcontroller)
anc_mc (ANC BC Microcontroller)
bc_bios (BC Tolapai BIOS)
lod_fpga (LOD FPGA)
node_bios (Node BIOS)
loc_fpga (LOC FPGA)
qloc_fpga (QLOC FPGA)
```

If the output of the `xtzap` command shows that only a specific type needs to be updated, use the `-t` option with that type (this example uses the `node_bios` type).

```
crayadm@smw:~> xtzap -t node_bios s0
```

If the output of the `xtzap` command shows that only blade component types need to be updated, use the `-b` option:

```
crayadm@smw:~> xtzap -b s0
```

If the output of the `xtzap` command shows that both blade- and cabinet-level component types need to be updated, or if you are unsure what needs to be updated, use the `-a` option:

```
crayadm@smw:~> xtzap -a s0
```

- d. Execute the `xtzap -r -v s0` command again; all firmware revisions should report correctly, except `node_bios`, which will display as "NOT_FOUND" until after the `xtbounce --linktune` command completes.

```
crayadm@smw:~> xtzap -r -v s0
```

- e. Execute the `xtbounce --linktune` command, which forces `xtbounce` to do full tuning on the system.

```
crayadm@smw:~> xtbounce --linktune=all s0
```

- f. Execute the `xtzap -r -v s0` command again, to verify that the BIOS version is now correct.

```
crayadm@smw:~> xtzap -r -v s0
```

5. As user `root`, place the Cray SMW 7.2UPxx Software DVD in the drive and mount it again.

```
smw:~ # mount /dev/cdrom /media/cdrom
```

6. Execute the `SMWconfig` command to configure MySQL. When you are prompted for the old root MySQL database password, press the Enter key (that is, an empty string), create a new root MySQL database password.

IMPORTANT: If you are unfamiliar with the requirements for setting and using the root MySQL password, in particular with regard to using the special characters `* ? [< > & ; ! | $`, review the MySQL documentation at oracle.com.

The SMWconfig command can be rerun with no adverse effects. For more detailed information about the SMWconfig command, see the SMWconfig(8) man page.

```
smw:~ # /media/cdrom/SMWconfig
17:00:33 Date started: Tue Oct 23 17:00:33 2012
17:00:33 Command Line used: /media/cdrom/SMWconfig
17:00:33 Validating command line options
. . .
Please enter your old root MySQL password:
Please confirm your old root MySQL password:
Password confirmed.
Please set your new root MySQL password:
Please confirm your new root MySQL password:
Password confirmed.
17:00:41 chkconfig rsms on
17:00:41 Checking the status of managers.
17:00:41 Setting MySQL root password.
17:00:41 Beginning to write the external mysql expect script.
17:00:41 Finished writing the external mysql expect script.
17:00:44 Beginning clean up...
17:00:44 Finished cleaning up.
17:00:44 SMWconfig has completed.
```

7. (If you do not configure postfix on your SMW skip this step.) To prevent your master.cf and main.cf postfix configuration files from being recreated during software updates or fixes, change the following setting in your /etc/sysconfig/mail file on your SMW to:

```
MAIL_CREATE_CONFIG="no"
```

8. Unmount the Cray SMW 7.2UPxx Software DVD.

```
smw:~ # umount /media/cdrom
smw:~ # exit
```

Confirm the SMW is Communicating with System Hardware

The following steps verify that the SMW is functional.

IMPORTANT: For a system configured for SMW high availability (HA) with the SMW failover feature, perform this procedure for the first SMW only. Skip this procedure for the second SMW.

1. Execute the xthwinv command on the entire system (s0) to examine the hardware inventory and verify that all nodes are visible to the SMW.

```
crayadm@smw:~> xthwinv s0
Received 24 of 24 responses.
Total number of reported modules: 24
Total number of reported nodes: 86
.
.
.
```

2. Execute the xtcli status on the entire system (s0).

```
crayadm@smw:~> xtcli status s0
Network topology: class 0
Network type: Gemini
Nodeid: Service Core Arch| Comp state [Flags]
```

```
-----
c0-0c0s0n0: service OP| on [noflags|]
c0-0c0s0n1: service OP| on [noflags|]
c0-0c0s0n2: service OP| on [noflags|]
c0-0c0s0n3: service OP| on [noflags|]
c0-0c0s1n0: - OP| on [noflags|]
c0-0c0s1n1: - OP| on [noflags|]
c0-0c0s1n2: - OP| on [noflags|]
c0-0c0s1n3: - OP| on [noflags|]
c0-0c0s2n0: service OP| on [noflags|]
c0-0c0s2n1: service OP| on [noflags|]
c0-0c0s2n2: service OP| on [noflags|]
c0-0c0s2n3: service OP| on [noflags|]
c0-0c0s3n0: - OP| on [noflags|]
c0-0c0s3n1: - OP| on [noflags|]
c0-0c0s3n2: - OP| on [noflags|]
c0-0c0s3n3: - OP| on [noflags|]
```

```
.
.
.
-----
```

3. Execute the `rtr -R` commands on the entire system (`s0`) to ensure that the SMW is functional. Note that the `rtr -R` command produces no output unless there is a routing problem.

```
crayadm@smw:~> rtr -R s0
```

4. Execute the `xtmcinfo -t -u` command to retrieve microcontroller information from cabinet control processors and blade control processors; for example:

```
crayadm@smw:~> xtmcinfo -t -u s0
c0-0c0s0 - Wed Apr 20 01:27:27 CDT 2012
c0-0c0s1 - Wed Apr 20 01:27:27 CDT 2012
c0-0c0s2 - Wed Apr 20 01:27:27 CDT 2012
c0-0c0s3 - Wed Apr 20 01:27:27 CDT 2012
c0-0c0s4 - Wed Apr 20 01:27:27 CDT 2012
c0-0c0s5 - Wed Apr 20 01:27:27 CDT 2012
c0-0c0s6 - Wed Apr 20 01:27:27 CDT 2012
c0-0c0s7 - Wed Apr 20 01:27:27 CDT 2012
.
.
.
c0-0c2s6 - Wed Apr 20 01:27:27 CDT 2012
c0-0c2s7 - Wed Apr 20 01:27:27 CDT 2012
How long have they been up...
c0-0c0s0 - 14:11:49 up 23 min, load average: 0.00, 0.01, 0.03
c0-0c0s1 - 14:11:49 up 23 min, load average: 0.53, 0.13, 0.06
c0-0c0s2 - 14:11:50 up 23 min, load average: 0.00, 0.00, 0.02
c0-0c0s3 - 14:11:50 up 23 min, load average: 0.14, 0.03, 0.01
c0-0c0s4 - 14:11:50 up 23 min, load average: 0.00, 0.01, 0.03
c0-0c0s5 - 14:11:50 up 23 min, load average: 0.01, 0.10, 0.12
c0-0c0s6 - 14:11:50 up 23 min, load average: 0.04, 0.10, 0.10
c0-0c0s7 - 14:11:50 up 23 min, load average: 0.01, 0.07, 0.08
.
.
```

```
.
c0-0c2s6 - 14:11:52 up 23 min, load average: 0.03, 0.11, 0.09
c0-0c2s7 - 14:11:53 up 23 min, load average: 0.01, 0.09, 0.08
```

Change the Default SMW Passwords

After completing the installation, change the default SMW passwords. The SMW contains its own `/etc/passwd` file that is separate from the password file for the rest of the system. To change the passwords on the SMW, log on to the SMW as `root` and execute the following commands:

```
crayadm@smw:~> su - root
smw:~# passwd root
smw:~# passwd crayadm
smw:~# passwd cray-vnc
smw:~# passwd mysql
```

To change the default iDRAC6 password see [Change the default iDRAC Password](#) on page 105.

Set up the SUSE firewall and IP tables

The SMW software includes a firewall. The following steps are for enabling and configuring the firewall on your system.

NOTE: You do not have to shut down the system to perform this task.

1. Before modifying the SUSE Firewall settings, make a copy of the configuration file:

```
smw:~ # cp -p /etc/sysconfig/SuSEfirewall12 /etc/sysconfig/SuSEfirewall12.orig
```

2. By using the `SuSEfirewall12` program and the following procedure, change the IP tables rules to close off all unnecessary ports on the SMW.

```
smw:~ # iptables -L
smw:~ # vi /etc/sysconfig/SuSEfirewall12
```

Change the settings of these variables to the values shown:

```
FW_DEV_EXT="any eth0"

FW_DEV_INT="eth1 eth2 eth3 eth4 lo"

FW_SERVICES_EXT_UDP="161"

FW_TRUSTED_NETS="your_bootnode_ipaddress,tcp,7004 \
                your_syslognode_ip,udp,514 your_sdbnode_ip,tcp,6811:6815"
```

For example:

```
smw:~ # diff /etc/sysconfig/SuSEfirewall12.orig /etc/sysconfig/SuSEfirewall12
99c99
< FW_DEV_EXT="eth-id-00:30:48:5c:b0:ee eth0"
```

```

---
> FW_DEV_EXT="any eth0"
114c114
< FW_DEV_INT="eth-id-00:0e:0c:b4:df:64 eth-id-00:0e:0c:b4:df:65
eth-id-00:0e:0c:b4:df:66 eth-id-00:0e:0c:b4:df:67 eth1 eth2 eth3 eth4"
---
> FW_DEV_INT="eth1 eth2 eth3 eth4 lo"
263c263
< FW_SERVICES_EXT_UDP=""
---
> FW_SERVICES_EXT_UDP="161"
394c394
< FW_TRUSTED_NETS=""
---
> FW_TRUSTED_NETS="10.3.1.254,tcp,7004 10.5.1.2,udp,514 10.5.1.2,tcp,6811:6815"

```

3. Invoke the modified configuration.

```

smw:~ # /etc/init.d/SuSEfirewall2_init start
smw:~ # /etc/init.d/SuSEfirewall2_setup start

```

4. Execute the following commands to start the firewall at boot time.

```

smw:~ # chkconfig SuSEfirewall2_init on
smw:~ # chkconfig SuSEfirewall2_setup on

```

5. To verify your changes to the iptables:

```

smw:~ # iptables -nvL

```

SSH access is one of the protocols permitted through the firewall from the external network to the SMW. For information about how to use Virtual Network Computing (VNC) through an SSH tunnel, see [Enable Remote Access to the SMW](#) on page 97.

Back Up Your Newly-installed SMW Software

After you have installed and configured your SMW software, create a backup of it.

- If you have a rack-mount SMW, see [Rack-mount SMW: Create a Cray System Management Workstation \(SMW\) Bootable Backup Drive](#).
- If you have a desk-side SMW, see [Desk-side SMW: Create an SMW Bootable Backup Drive](#) on page 43.

Rack-mount SMW: Create an SMW bootable backup drive

IMPORTANT: The disk device names shown in this procedure are only examples. You should substitute the actual disk device names for your system. The boot disk is phy7 and is slot 0, and the bootable backup disk is phy6 and is slot 1.



CAUTION: Shut down the Cray system before you begin this procedure.

Also be aware that there may be a considerable load on the SMW while creating the SMW bootable backup drive.

1. Log on to the SMW as `crayadm` and `su` to `root`.

```
crayadm@smw:~> su -
Password:
smw:~ # ls -al /dev/disk/by-path
total 0
drwxr-xr-x 2 root root 380 Mar 15 13:21 .
drwxr-xr-x 6 root root 120 Mar 11 18:42 ..
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:00:11.0-scsi-0:0:0:0 -> ../../sr0
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:00:12.2-usb-0:3:1.0-scsi-0:0:0:0 -
> ../../sdf
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:00:12.2-usb-0:3:1.1-scsi-0:0:0:0 -
> ../../sr1
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:00:12.2-usb-0:3:1.1-scsi-0:0:0:1 -
> ../../sdg
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:00:13.2-usb-0:2:1:1.0-scsi-0:0:0:0 -
> ../../sde
lrwxrwxrwx 1 root root 10 Mar 11 18:42 pci-0000:00:13.2-usb-0:2:1:1.0-scsi-0:0:0:0-
part1 -> ../../sde1
lrwxrwxrwx 1 root root 10 Mar 11 18:42 pci-0000:00:13.2-usb-0:2:1:1.0-scsi-0:0:0:0-
part2 -> ../../sde2
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:05:00.0-sas-phy4-...
lrwxrwxrwx 1 root root 10 Mar 14 15:57 pci-0000:05:00.0-sas-phy4-...
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:05:00.0-sas-phy5-...
lrwxrwxrwx 1 root root 10 Mar 14 16:00 pci-0000:05:00.0-sas-phy5-...
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:05:00.0-sas-phy6-...
lrwxrwxrwx 1 root root 10 Mar 15 13:21 pci-0000:05:00.0-sas-phy6-...
lrwxrwxrwx 1 root root 10 Mar 15 13:21 pci-0000:05:00.0-sas-phy6-...
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:05:00.0-sas-phy7-...
lrwxrwxrwx 1 root root 10 Mar 11 18:42 pci-0000:05:00.0-sas-phy7-...
lrwxrwxrwx 1 root root 10 Mar 11 18:42 pci-0000:05:00.0-sas-phy7-...
```

2. Standardize the SMW's boot-time drive names with the Linux run-time drive names.

IMPORTANT: If the SMW configuration files on the SMW root drive have been modified already (because your site has completed this step at least once after installing your updated SMW base operating system), skip to [3](#) on page 37 ; otherwise, complete this step to standardize the SMW's boot-time drive names with the Linux run-time drive names.

Set up ordered drives on your rack-mount SMW.

- a. Identify the installed SMW drive model numbers, [3](#) on page 37 serial numbers, and associated Linux device (`/dev`) names.

Execute the `smwmapdrives` script on the SMW to identify local (internal) drives mounted in the SMW and provide their Linux device (`/dev`) names.

NOTE: Effective with the SMW 7.2.UP00 release, the `smwmapdrives` script was provided both as a separate file in the release and in the base operating system RPM. If you are running that release or a later one, you can use the installed version of the script to back up the SMW.

```
smw:~ # ./smwmapdrives
List of SMW-installed disk drives
-----
Physical slot 0:
    /dev/sda
    /dev/disk/by-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RDS
    /dev/disk/by-id/scsi-SATA_FUJITSU_MHZ2160_K85DTB227RDS
    /dev/disk/by-path/pci-0000:05:00.0-sas-phy7-0x4433221107000000-lun-0
Physical slot 1:
    /dev/sdc
    /dev/disk/by-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RD7
    /dev/disk/by-id/scsi-SATA_FUJITSU_MHZ2160_K85DTB227RD7
    /dev/disk/by-path/pci-0000:05:00.0-sas-phy6-0x4433221106000000-lun-0
Physical slot 2:
```

```

/dev/sdd
/dev/disk/by-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RF3
/dev/disk/by-id/scsi-SATA_FUJITSU_MHZ2160_K85DTB227RF3
/dev/disk/by-path/pci-0000:05:00.0-sas-phy5-0x4433221105000000-lun-0
Physical slot 3:
/dev/sdb
/dev/disk/by-id/ata-ST9500620NS_9XF0665V
/dev/disk/by-id/scsi-SATA_ST9500620NS_9XF0665V
/dev/disk/by-path/pci-0000:05:00.0-sas-phy4-0x4433221104000000-lun-0
Physical slot 4:
NOT INSTALLED
Physical slot 5:
NOT INSTALLED

```

The device names for `by-id` are persistent and will reference the drive, regardless of the slot in which the drive is installed.

`by-path` names reference a physical drive slot only and do not identify the drive installed in that slot. This is the naming used by default for the logging and database drives when the SMW was installed. This `by-path` name is used to specifically install logging and database file systems because the `by-id` device names refer to the physical drive slots expected to be used for those file systems and are provided as the default examples in the SMW installation configuration process.

The `/dev/sdX` drive names are not persistent, except for the `/dev/sda` drive name, which always refers to the drive installed in physical slot 0 (zero) of the SMW. The other `/dev/sdX` names can change with each SMW boot and will change if drives are added, removed, or reordered in the SMW slots. For this reason, the `/dev/sda` drive name can only be used for the rack-mount SMW.

Choose either the `by-id` naming or the `by-path` naming as the site administrative policy for managing the SMW-install disk drives. The following documentation provides the steps necessary to implement this selection on the SMW prior to creating an SMW bootable backup drive.

- b. Back up the following files before proceeding:

```

smw# cp -p /boot/grub/device.map /boot/grub/device.map-YYYYMMDD
smw# cp -p /boot/grub/menu.lst /boot/grub/menu.lst-YYYYMMDD
smw# cp -p /etc/fstab /etc/fstab-YYYYMMDD

```

- c. Edit the grub `device.map` file to reflect physical drive locations.

To provide a direct mapping of the SMW disk drive physical slots to the boot loader (BIOS and grub) drive names, the `device.map` mapping file used by grub should be replaced. Perform the following steps to install new `device.map` file entries to effect this mapping.

1. Edit the grub `device.map` file.
2. Delete all lines.
3. Enter the following lines into the file. These lines show each drive slot's physical location mapped to its boot-time `hd?` name.

NOTE: `by-id` names should not be used in the `device.map` file.

```

# Dell Rackmount r805 SMW
# grub(8) device mapping for boot-drive identification
# hd? numbers are being mapped to their physical
(hd0) /dev/disk/by-path/pci-0000:05:00.0-sas-
phy7-0x4433221107000000-lun-0
(hd1) /dev/disk/by-path/pci-0000:05:00.0-sas-
phy6-0x4433221106000000-lun-0
(hd2) /dev/disk/by-path/pci-0000:05:00.0-sas-
phy5-0x4433221105000000-lun-0

```

```
(hd3)    /dev/disk/by-path/pci-0000:05:00.0-sas-
phy4-0x4433221104000000-lun-0
(hd4)    /dev/disk/by-path/pci-0000:05:00.0-sas-
phy3-0x4433221103000000-lun-0
(hd5)    /dev/disk/by-path/pci-0000:05:00.0-sas-
phy2-0x4433221102000000-lun-0
```

- d. Modify the SMW boot drive `/etc/fstab` file to use `by-id` or `by-path` naming.

Modify the SMW file system mounting configuration file to use SMW disk `by-id` or `by-path` naming. Complete this step to replace any `/dev/sdX` disk partition references.

NOTE: Use the output of the `smwmapdrives` script in as a reference for drive names.

Edit `/etc/fstab`, replacing drive `/dev/sdX` references with either the `by-id` or `by-path` name's corresponding device name.

When a reference to `/dev/sda1` is being replaced, replace it with the corresponding "partition" file system suffixed with `-part1`. File system partitions for `/dev/sda` are indicated by the numeral appended to the device name; for example, `/dev/sda1` refers to partition 1 on `/dev/sda`. Replace it with the `by-id` and `by-path` device names and suffix the device name with `-part1`, keeping the same numeral identification.

For example, if the root and swap file systems are currently configured to mount `/dev/sda2`, they should be changed. Using the `by-id` device name from the example in , the `fstab` lines would change from:

```
/dev/sda1 swap                swap          defaults      0 0
/dev/sda2 /                    ext3          acl,user_xattr 1 1to:
/dev/disk/by-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RDS-part1 swap
swap defaults                0 0
/dev/disk/by-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RDS-part2 /
ext3 acl,user_xattr          1 1
```

- e. Modify `/boot/grub/menu.lst` to reflect the `device.map` BIOS/boot-up drive changes for the `sdX` remapping.

The same device name replacement performed on `/etc/fstab` should also be performed on the grub bootloader `/boot/grub/menu.lst` configuration file. All references to `/dev/sdX` devices should be replaced with corresponding `by-id` or `by-path` device names.

- f. Invoke the grub utility to reinstall the SMW boot loader on the primary boot drive.

Once the changes to `device.map`, `fstab`, and `menu.lst` have been completed, the grub bootloader boot blocks must be updated to reflect changes to the device names. Complete this step to update the boot loader on the boot drive.

Invoke the grub utility and reinstall SMW root-drive boot blocks.

```
smw:~ # grub --no-curses
GNU GRUB version 0.97 (640K lower / 3072K upper memory)
[ Minimal BASH-like line editing is supported. For the first word, TAB
  lists possible command completions. Anywhere else TAB lists the possible
  completions of a device/filename. ]
grub> root (hd0,1)
      root (hd0,1)
      Filesystem type is ext2fs, partition type 0x83
grub> setup (hd0)
      Checking if "/boot/grub/stage1" exists... yes
```

```

Checking if "/boot/grub/stage2" exists... yes
Checking if "/boot/grub/e2fs_stage1_5" exists... yes
Running "embed /boot/grub/e2fs_stage1_5 (hd0)"... 17 sectors \
are embedded. Succeeded
Running "install /boot/grub/stage1 (hd0) (hd0)1+17 p (hd0,1)/boot/grub/
stage2 \
/boot/grub/menu.lst"... succeeded
Done.
grub> quit

```

3. If the backup drive disk partition table already exists and the partition table on the backup drive matches the partition table that is on the primary boot drive, skip this step; otherwise, create the backup drive disk partition table.

In this example, the partition table consists of two slices. Slice 1 is a 4 GB Linux swap partition. Slice 2 is the balance of disk space used for the root file system.

- a. Use the `fdisk` command to display the boot disk partition layout.

```

smw:~ # fdisk -lu \
/dev/disk/by-path/pci-0000:05:00.0-sas-phy7-0x4433221107000000-lun-0
Disk /dev/disk/by-path/pci-0000:05:00.0-sas-phy7-0x4433221107000000-lun-0: 250.0 GB, \
268435456000 bytes
255 heads, 63 sectors/track, 19457 cylinders, total 312581808 sectors
Units = sectors of 1 * 512 = 512 bytes
Disk identifier: 0x00000082

    Device
/dev/disk/by-path/pci-0000:05:00.0-sas-phy7-0x4433221107000000-lun-0-part1 \
    Boot   Start       End     Blocks  Id System
          63   16771859   8385898+  82  Linux swap / Solaris
/dev/disk/by-path/pci-0000:05:00.0-sas-phy7-0x4433221107000000-lun-0-part2 \
    Boot   Start       End     Blocks  Id System
          * 16771860 312576704 147902422+ 83  Linux

```

- b. Use the `fdisk` command to configure the bootable backup disk partition layout. Set the bootable backup disk partition layout to match the boot disk partition layout. First, clear all of the old partitions using the `d` command within `fdisk`; next create a Linux swap and a Linux partition; and then write your changes to the disk. For help, type `m` within `fdisk`.

```

smw:~ # fdisk -u \
/dev/disk/by-path/pci-0000:05:00.0-sas-phy6-0x4433221106000000-lun-0

The number of cylinders for this disk is set to 19457.
There is nothing wrong with that, but this is larger than 1024,
and could in certain setups cause problems with:
1) software that runs at boot time (e.g., old versions of LILO)
2) booting and partitioning software from other OSs
   (e.g., DOS FDISK, OS/2 FDISK)

Command (m for help): p

Disk /dev/disk/by-path/pci-0000:05:00.0-sas-phy6-0x4433221106000000-lun-0: 250.0 GB, \
268435456000 bytes
255 heads, 63 sectors/track, 19457 cylinders, total 312581808 sectors
Units = sectors of 1 * 512 = 512 bytes
Disk identifier: 0x00000080

    Device
/dev/disk/by-path/pci-0000:05:00.0-sas-phy6-0x4433221106000000-lun-0-part1 \
    Boot   Start       End     Blocks  Id System
          63   16771859   83828    82  Linux
swap / Solaris

```

Partition 1 does not end on cylinder boundary.

```
/dev/disk/by-path/pci-0000:05:00.0-sas-phy6-0x4433221106000000-lun-0-part2 \
      Boot   Start       End     Blocks  Id  System
      167719 312581807 156207044+  83  Linux
```

```
Command (m for help): d
Partition number (1-4): 2
```

```
Command (m for help): d
Selected partition 1
```

```
Command (m for help): n
Command action
  e   extended
  p   primary partition (1-4)
```

```
p
Partition number (1-4): 1
First sector (63-312581807, default 63): (Press the Enter key)
Using default value 63
Last sector, +sectors or +size{K,M,G} (63-312581807, default 312581807): 16771859
Command (m for help): t
Selected partition 1
Hex code (type L to list codes): 82
Changed system type of partition 1 to 82 (Linux swap / Solaris)
```

```
Command (m for help): n
Command action
  e   extended
  p   primary partition (1-4)
```

```
p
Partition number (1-4): 2
First sector (16771860-312581807, default 16771860): (Press the Enter key)
Using default value 16771860
Last sector, +sectors or +size{K,M,G} (16771860-312581807, default 312581807): (Press
the Enter key)
Using default value 312581807
```

```
Command (m for help): w
The partition table has been altered!
```

```
Calling ioctl() to re-read partition table.
Syncing disks.
```

- c. Display the boot backup disk partition layout and confirm it matches your phy7 sector information.

```
smw:~ # fdisk -lu \
/dev/disk/by-path/pci-0000:05:00.0-sas-phy6-0x4433221106000000-lun-0
Disk /dev/disk/by-path/pci-0000:05:00.0-sas-phy6:1-0x4433221106000000:0-lun0: 250.0
GB, \
268435456000 bytes
255 heads, 63 sectors/track, 19457 cylinders, total 312581808 sectors
```

4. Initialize the swap device.

```
smw:~ # mkswap \
/dev/disk/by-path/pci-0000:05:00.0-sas-phy6-0x4433221106000000-lun-0-part1
mkswap: /dev/disk/by-path/pci-0000:05:00.0-sas-phy6:1-0x4433221106000000:0-lun0-part1:
warning: don't erase bootbits sectors
      (DOS partition table detected). Use -f to force.
Setting up swapspace version 1, size = 8385892 KiB
no label, UUID=c0ef22ac-b405-4236-855b-e4a09b6e94ed
```

5. Create a new file system on the backup drive root partition by executing the mkfs command.

```
smw:~ # mkfs -t ext3 \
/dev/disk/by-path/pci-0000:05:00.0-sas-phy6-0x4433221106000000-lun-0-part2
```

```

mke2fs 1.41.9 (22-Aug-2009)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
9248768 inodes, 36976243 blocks
1848812 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=4294967296
1129 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000, 7962624, 11239424, 20480000, 23887872

Writing inode tables:   done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 37 mounts or
180 days, whichever comes first.  Use tune2fs -c or -i to override.

```

6. Mount the new backup root file system on /mnt.

```

smw:~ # mount \
/dev/disk/by-path/pci-0000:05:00.0-sas-phy6-0x4433221106000000-lun-0-part2 /mnt

```

7. Confirm that the backup root file system is mounted.

```

smw:~ # df
Filesystem            1K-blocks      Used Available Use% Mounted on
/dev/sda2              303528624    6438700 281671544   3% /
udev                  1030332        116   1030216    1% /dev
/dev/sdb2             306128812    195568 290505224   1% /mnt

```

The running root file system device is the one mounted on /.

8. Dump the running root file system to the backup drive.

```

smw:~ # cd /mnt
smw:~ # dump 0f - \
/dev/disk/by-path/pci-0000:05:00.0-sas-phy7-0x4433221107000000-lun-0-part2 | restore rf -
DUMP: WARNING: no file `/etc/dumpdates'
DUMP: Date of this level 0 dump: Tue Mar 15 13:43:17 2011
DUMP: Dumping /dev/sda2 (/) to standard output
DUMP: Label: none
DUMP: Writing 10 Kilobyte records
DUMP: mapping (Pass I) [regular files]
DUMP: mapping (Pass II) [directories]
DUMP: estimated 7898711 blocks.
DUMP: Volume 1 started with block 1 at: Tue Mar 15 13:44:40 2011
DUMP: dumping (Pass III) [directories]
DUMP: dumping (Pass IV) [regular files]
restore: ./lost+found: File exists
DUMP: 79.34% done at 20890 kB/s, finished in 0:01
DUMP: Volume 1 completed at: Tue Mar 15 13:52:13 2011
DUMP: Volume 1 7908080 blocks (7722.73MB)
DUMP: Volume 1 took 0:07:33
DUMP: Volume 1 transfer rate: 17457 kB/s
DUMP: 7908080 blocks (7722.73MB)
DUMP: finished in 453 seconds, throughput 17457 kBytes/sec
DUMP: Date of this level 0 dump: Tue Mar 15 13:43:17 2011
DUMP: Date this dump completed: Tue Mar 15 13:52:13 2011

```

```
DUMP: Average transfer rate: 17457 kB/s
DUMP: DUMP IS DONE
```

9. Modify the backup drive's `fstab` and `menu.lst` files to reflect the backup drive's device, replacing the primary drive's device name.

NOTE: This step is necessary only if `by-id` names are used. If `by-path` names are being utilized for the root and swap devices, changes are not necessary; these devices reference physical slots, and the backup drive will be moved to the same physical slot (slot 0) when replacing a failed primary boot drive.

- a. Edit `/mnt/etc/fstab`. Replace the root and swap partitions' `by-id` device names with those used for this backup device, replacing the original disk device name.

For example, change

```
/dev/disk/by-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RDS-part1 swap swap defaults
/dev/disk/by-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RDS-part2 / ext3 acl,user_xattr
```

to:

```
/dev/disk/by-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RD7-part1 swap swap defaults
/dev/disk/by-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RD7-part2 / ext3 acl,user_xattr
```

- b. Edit `/mnt/boot/grub/menu.lst`. Replace the `root=` and `resume=` device names with those used for this backup device, replacing the original disk device name.

The `root=` entry normally refers to partition `-part2`, and the `resume=` entry normally refers to partition `-part1`; these partition references must be maintained.

For example, replace the `menu.lst` configuration references of:

```
root=/dev/disk/by-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RDS-part2
```

with:

```
root=/dev/disk/by-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RD7-part2or
similarly with the by-id device names, if those are preferred.
```

Replace the `resume=` references similarly.

10. Install the `grub` boot loader. To make the backup drive bootable, reinstall the `grub` boot facility on that drive.



CAUTION: Although all of the disks connected to the SMW are available to the system, `grub` only detects the first 16 devices. Therefore, if you add a disk to the SMW after the SMW is connected to the boot RAID, it is advisable to reboot the SMW before continuing.

- a. Create a unique file on the backup drive to be used to identify that drive to `grub` boot facility.

```
smw:~ # cd /
smw:~ # touch /mnt/THIS_IS_6
```

- b. Invoke the `grub` boot utility. Within the `grub` boot utility:

1. Execute the `find` command to locate the drive designation that `grub` uses.
2. Select the drive to which the boot blocks will be installed with the `root` command.

3. Use the setup command to set up and install the grub boot blocks on that drive. The Linux grub utility and boot system always refer to drives as `hd`, regardless of the actual type of drives. For example:

```
smw:~ # grub --no-curses
GNU GRUB version 0.97 (640K lower / 3072K upper memory)
[ Minimal BASH-like line editing is supported. For the first word, TAB
  lists possible command completions. Anywhere else TAB lists the
possible
  completions of a device/filename. ]
grub> find /THIS_IS_6
(hd2,1)
grub> root (hd2,1)
root (hd2,1)
Filesystem type is ext2fs, partition type 0x83
grub> setup (hd2)
Checking if "/boot/grub/stage1" exists... yes
Checking if "/boot/grub/stage2" exists... yes
Checking if "/boot/grub/e2fs_stage1_5" exists... yes
Running "embed /boot/grub/e2fs_stage1_5 (hd2)"... 17 sectors are
embedded.
succeeded
Running "install /boot/grub/stage1 (hd2) (hd2)1+17 p (hd2,1)/boot/grub/
stage2 \
/boot/grub/menu.lst"... succeeded
Done.
grub> quit
```

IMPORTANT: For rack-mount SMWs, grub recreates `device.map` with the short names, not the persistent names. The `/dev/sdx` names must not be trusted. Always use `find` the next time you execute grub because it is possible that grub `root` may not be `hd2` the next time you execute grub.

11. Unmount the backup root partition.

```
smw:~ # umount /mnt
```

The drive is now bootable once plugged in or cabled as the primary drive.

Rack-mount SMW: Set Up the Bootable Backup Drive as an Alternate Boot Device

The following procedure modifies a bootable backup drive for a rack-mount SMW in order to boot from and run the rack-mount SMW from the backup root partition.

IMPORTANT: In order to boot from this backup drive, the primary boot drive must still be operable and able to boot the grub boot blocks installed. If the backup drive is modified to boot as an alternate boot device, it will no longer function as a bootable backup if the primary drive fails.

The disk device names shown in this procedure are only examples. You should substitute the actual disk device names for your system. The boot disk is `phy7` and is slot 0, and the bootable backup disk is `phy6` and is slot 1

1. Mount the backup drive's root partition.

```
smw:~ # mount /dev/disk/by-path/pci-0000:05:00.0-sas-phy6-0x4433221106000000-
lun-0-part2 /mnt
```

2. Create a new boot entry in the `/boot/grub/menu.lst` file. This entry should be a duplicate of the primary boot entry with the following changes:
 - a. Modify the title to uniquely identify the backup boot entry.
 - b. Modify the `root (hd0,1)` directive to reflect the grub name of the backup drive.
 - c. Modify the `root=` and `resume=` specifications to reference the backup drive device.

This is an example `/boot/grub/menu.lst` file. Note the new entry for the backup drive. This example references `phy7` (slot 0) and as the primary drive and `phy6` (slot 1) as the backup drive.

```
smw:~ # cp -p /boot/grub/menu.lst /boot/grub/menu.lst.20110317
smw:~ # vi /boot/grub/menu.lst
smw:~ # cat /boot/grub/menu.lst
# Modified by YaST2. Last modification on Wed Jun 27 12:32:43 CDT 2012
default 0
timeout 8
##YaST - generic_mbr
gfxmenu (hd0,1)/boot/message
##YaST - activate

###Don't change this comment - YaST2 identifier: Original name: linux###
title SUSE Linux Enterprise Server 11 SP3 - 3.0.26-0.7
    root (hd0,1)
    kernel /boot/vmlinuz-3.0.26-0.7-default \
    root=/dev/disk/by-path/pci-0000:05:00.0-sas-phy7-0x4433221107000000-lun-0-
part2 \
    resume=/dev/disk/by-path/pci-0000:05:00.0-sas-phy7-0x4433221107000000-lun-0-
part1 \
    splash=silent crashkernel=256M-:128M showopts vga=0x31a
    initrd /boot/initrd-3.0.26-0.7-default

### New entry allowing a boot of the back-up drive when the primary drive
### is still present.
title BACK-UP DRIVE - SUSE Linux Enterprise Server 11 SP3 - 3.0.26-0.7
    root (hd1,1)
    kernel /boot/vmlinuz-3.0.26-0.7-default \
    root=/dev/disk/by-path/pci-0000:05:00.0-sas-phy6-0x4433221106000000-lun-0-
part2 \
    resume=/dev/disk/by-path/pci-0000:05:00.0-sas-phy6-0x4433221106000000-lun-0-
part1 \
    splash=silent crashkernel=256M-:128M showopts vga=0x31a
    initrd (hd0,1)/boot/initrd-3.0.26-0.7-default

###Don't change this comment - YaST2 identifier: Original name: failsafe###
title Failsafe -- SUSE Linux Enterprise Server 11 SP3 - 3.0.26-0.7
    root (hd0,1)
    kernel /boot/vmlinuz-3.0.26-0.7-default \
    root=/dev/disk/by-path/pci-0000:05:00.0-sas-phy7-0x4433221107000000-lun-0-
part2 \
    showopts ide=nodma apm=off noresume edd=off powersaved=off \
    nohz=off highres=off processor.max_cstate=1 nomodeset x11failsafe vga=0x31a
    initrd /boot/initrd-3.0.26-0.7-default
```

3. Modify the backup drive's `/etc/fstab` file to reference the secondary drive slot rather than the first drive slot.
 - a. Examine the backup drive's `fstab` file.

```

smw:~ # cp -p /mnt/etc/fstab /mnt/etc/fstab.20110317
smw:~ # cat /mnt/etc/fstab
/dev/disk/by-path/pci-0000:05:00.0-sas-phy7-0x4433221107000000-lun-0-part1 \
swap      swap      defaults      0 0
/dev/disk/by-path/pci-0000:05:00.0-sas-phy7-0x4433221107000000-lun-0-part2 \
/          ext3      acl,user_xattr 1 1
proc       /proc           proc           defaults      0 0
sysfs      /sys           sysfs          noauto        0 0
debugfs    /sys/kernel/debug debugfs        noauto        0 0
usbfs      /proc/bus/usb  usbfs          noauto        0 0
devpts     /dev/pts       devpts         mode=0620,gid=5 0 0

```

- b. Edit the /mnt/etc/fstab file, changing phy7 to phy6 device names to reference the backup drive. In the following example, the backup drive is phy6:1-....

```

smw:~ # vi /mnt/etc/fstab
smw:~ # cat /mnt/etc/fstab
/dev/disk/by-path/pci-0000:05:00.0-sas-phy6-0x4433221106000000-lun-0-part1 \
swap      swap      defaults      0 0
/dev/disk/by-path/pci-0000:05:00.0-sas-phy6-0x4433221106000000-lun-0-part2 \
/          ext3      acl,user_xattr 1 1
proc       /proc           proc           defaults      0
0
sysfs      /sys           sysfs          noauto        0
0
debugfs    /sys/kernel/debug debugfs        noauto        0
0
usbfs      /proc/bus/usb  usbfs          noauto        0
0
devpts     /dev/pts       devpts         mode=0620,gid=5 0
0

```

4. Unmount the backup drive.

```
smw:~ # umount /mnt
```

The SMW can now be shut down and rebooted. Upon display of the **Please select boot device** prompt, select the BACK-UP DRIVE - SLES 11 entry to boot the backup root partition.

Desk-side SMW: Create an SMW Bootable Backup Drive

IMPORTANT: The disk device names shown in this procedure are only examples; substitute the actual disk device names for your system. For example, on an SMW with three SMW disks, the boot disk is /dev/sda and the bootable backup disk is /dev/sdc; on an SMW with two SMW disks, the boot disk is /dev/sda and the bootable backup disk is /dev/sdb.



CAUTION: Shut down the Cray system before you begin this procedure.

Also be aware that there may be a considerable load on the SMW while creating the SMW bootable backup drive.

1. Log on to the SMW as `crayadm` and `su` to `root`.

```

crayadm@smw:~> su - root
smw:~ #

```

2. If the backup drive disk partition table already exists and the partition table on the backup drive matches the partition table that is on the primary boot drive, skip this step; otherwise, create the backup drive disk partition table.

NOTE:

For optimal performance, the source and destination disks should be on different buses; drive slots 0 and 1 are on a different bus than drive slots 2 and 3.

In this example, the partition table consists of the following:

- Slice 1: 4 GB Linux swap partition
- Slice 2: Balance of disk space used for the root file system

- a. Use the `fdisk` command to display the boot disk partition layout.

```
smw:~ # fdisk -lu /dev/sda
Disk /dev/sda: 320.0 GB, 320072933376 bytes
255 heads, 63 sectors/track, 38913 cylinders, total 625142448 sectors
Units = sectors of 1 * 512 = 512 bytes
```

Device	Boot	Start	End	Blocks	Id	System
/dev/sda1		63	8401994	4200966	82	Linux swap / Solaris
/dev/sda2	*	8401995	625137344	308367675	83	Linux

- b. Use the `fdisk` command to configure the bootable backup disk partition layout. Set the bootable backup disk partition layout to match the boot disk partition layout. First, clear all of the old partitions using the `d` command within `fdisk`; next create a Linux swap and a Linux partition; and then write your changes to the disk. For help, type `m` within `fdisk` (see the following sample output).

```
smw:~ # fdisk -u /dev/sdb

The number of cylinders for this disk is set to 38913.
There is nothing wrong with that, but this is larger than 1024,
and could in certain setups cause problems with:
1) software that runs at boot time (e.g., old versions of LILO)
2) booting and partitioning software from other OSs
   (e.g., DOS FDISK, OS/2 FDISK).

Command (m for help): p
Disk /dev/sdb: 320.0 GB, 320072933376 bytes
255 heads, 63 sectors/track, 38913 cylinders, total 625142448 sectors
Units = sectors of 1 * 512 = 512 bytes
```

Device	Boot	Start	End	Blocks	Id	System
/dev/sdb1		63	8401994	4200966	82	Linux swap
/dev/sdb2		8401995	625105214	308351610	83	Linux

```
Command (m for help): d
Partition number (1-5): 2
Command (m for help): d
Selected partition 1
Command (m for help): n
Command action
e extended
p primary partition (1-4)
p
Partition number (1-4): 1
First sector (63-625105215, default 63): (Press the Enter key)
Using default value 63
```

```
Last sector or +size or +sizeM or +sizeK (63-625105215, default 625105215):
8401994
```

```
Command (m for help): t
Selected partition 1
Hex code (type L to list codes): 82
Changed system type of partition 1 to 82 (Linux swap / Solaris)

Command (m for help): n
Command action
e extended
p primary partition (1-4)
p
Partition number (1-4): 2
First sector (8401995-625105215, default 8401995): (Press the Enter key)
Using default value 8401995
Last sector or +size or +sizeM or +sizeK (8401995-625105215, default
625105215): \
(Press the Enter key)
Using default value 625105215

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
```

- c. Display the boot backup disk partition layout.

```
smw:~ # fdisk -lu /dev/sdb
Disk /dev/sdb: 320.0 GB, 320072933376 bytes
255 heads, 63 sectors/track, 38913 cylinders, total 625142448 sectors
Units = sectors of 1 * 512 = 512 bytes
```

Device	Boot	Start	End	Blocks	Id	System
/dev/sdc1		63	8401994	4200966	82	Linux swap / Solaris
/dev/sdc2	*	8401995	625137344	308367675	83	Linux

3. Initialize the swap device.

```
smw:~ # mkswap /dev/sdb1
```

4. (If the device names have already been standardized, skip this step.) Standardize the `/etc/fstab` and `grub` disk device names. The device names that the installation process writes into the `/boot/grub/menu.lst` file are UDEV-based names (for example, `/dev/disk/by-id/scsi-SATA_ST3320620AS_922J3-part2` or `/dev/disk/by-id/ata-ST3320620A_9QFA85PV-part2`) instead of the more commonly used device names (for example, `/dev/sda2` or `/dev/hda2`). In the following procedures, edit the `/boot/grub/menu.lst` file to change only the long UDEV-based name to the shorter, commonly used device name reflected in the output of the `df` command on your system. Be aware that errors in the `/boot/grub/menu.lst` will affect your ability to boot the SMW.

- a. SLES 11 sets up `/etc/fstab` and `/boot/grub/menu.lst` with UDEV-based names for the root device. For example:

```
smw:~ # head -2 /etc/fstab
/dev/disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part2 / ext3 acl,user_xattr 1
```

```

1
/dev/disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part1 swap swap defaults 0 0
smw:~ # more /boot/grub/menu.lst
###Don't change this comment - YaST2 identifier: Original name: linux###
title SUSE Linux Enterprise Server 11 - 2.6.27.19-5
root (hd0,1)
kernel /boot/vmlinuz-2.6.27.19-5-default \
root=/dev/disk/by-id/ata-ST3320620AS_5QF00F84-part2 \
resume=/dev/sda1 splash=silent crashkernel=256M-:128M@16M
showopts vga=0x31a initrd /boot/initrd-2.6.27.19-5-default
###Don't change this comment - YaST2 identifier: Original name: failsafe###
title Failsafe -- SUSE Linux Enterprise Server 11 - 2.6.27.19-5
root (hd0,1)
kernel /boot/vmlinuz-2.6.27.19-5-default \
root=/dev/disk/by-id/ata-ST3320620AS_5QF00F84-part2 showopts \
ide=nodma apm=off noresume edd=off powersaved=off nohz=off highres=off
processor.max_cstate=1 x11failsafe vga=0x31a
initrd /boot/initrd-2.6.27.19-5-default

```

- b. Execute the `df` command to get the name of the device to use in the `/etc/fstab` and `/boot/grub/menu.lst` files to replace the long UDEV-based device name.

For example:

```

smw:# df
Filesystem      1K-blocks    Used    Available    Use%    Mounted on
/dev/sda2        303528624 40652904   247457340    15%     /
udev             1030780    460       1030320      1%      /dev

```

- c. Create a backup copy of the `/etc/fstab` and `/boot/grub/menu.lst` files.

```

smw:# cp -p /etc/fstab /etc/fstab.save
smw:# cp -p /boot/grub/menu.lst /boot/grub/menu.lst.save

```

- d. Edit the `/etc/fstab` file appropriately, using the device name (`dev`) from the `df` command output. In this example, the "1" and "2" refer to the partition names on the device. Change the following lines, which changes the long name `disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part2` to `sda2` and changes `disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part1` to `sda1`. Ensure that your swap is on `sda1`:

```

/dev/disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part2 / ext3 acl,user_xattr 1
1
/dev/disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part1 swap swap defaults 0 0
to:
/dev/sda2 / ext3 acl,user_xattr 1 1
/dev/sda1 swap swap defaults 0 0

```

- e. Edit the `/boot/grub/menu.lst` file appropriately; use the device name (`dev`) from the `df` command output. Change the long name `disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part2` to `sda2`. Change the following lines:

```

title SUSE Linux Enterprise Server 11 - 2.6.27.19-5
kernel /boot/vmlinuz-2.6.27.19-5-default \
root=/dev/disk/by-id/ata-ST3320620AS_5QF00F84-part2 \
resume=/dev/sda1 splash=silent crashkernel=256M-:128M@16M showopts vga=0x31a
to

```

```
title SUSE Linux Enterprise Server 11 - 2.6.27.19-5
kernel /boot/vmlinuz-2.6.27.19-5-default \
root=/dev/sda2 resume=/dev/sda1 splash=silent \
crashkernel=256M-:128M@16M showopts vga=0x31a
```

and change the following lines:

```
title Failsafe -- SUSE Linux Enterprise Server 11 - 2.6.27.19-5
kernel /boot/vmlinuz-2.6.27.19-5-default \
root=/dev/disk/by-id/ata-ST3320620AS_5QF00F84part2 \
showopts ide=nodma apm=off noresume edd=off powersaved=off nohz=off \
highres=off processor.max_cstate=1 x11failsafe vga=0x31a
```

to

```
title Failsafe -- SUSE Linux Enterprise Server 11 - 2.6.27.19-5
kernel /boot/vmlinuz-2.6.27.19-5-default \
root=/dev/sda2 showopts ide=nodma apm=off noresume edd=off \
powersaved=off nohz=off highres=off processor.max_cstate=1 x11failsafe
vga=0x31a
```

- f. Verify that the edited files are correct and match the output of the `df` command.

```
smw:~ # head -2 /etc/fstab
/dev/sda2 / ext3 acl,user_xattr 1 1
/dev/sda1 swap swap defaults 0 0
smw:~ # more /boot/grub/menu.lst
###Don't change this comment - YaST2 identifier: Original name: linux###
title SUSE Linux Enterprise Server 11 - 2.6.27.19-5
root (hd0,1)
kernel /boot/vmlinuz-2.6.27.19-5-default root=/dev/sda2 \
resume=/dev/sda1 splash=silent crashkernel=256M-:128M@16M showopts
initrd /boot/initrd-2.6.27.19-5-default
###Don't change this comment - YaST2 identifier: Original name: failsafe###
title Failsafe -- SUSE Linux Enterprise Server 11 - 2.6.27.19-5
root (hd0,1)
kernel /boot/vmlinuz-2.6.27.19-5-default \
root=/dev/sda2 showopts ide=nodma apm=off noresume edd=off
powersaved=off nohz=off highres=off initrd /boot/initrd-2.6.27.19-5-default
```

5. Update the grub device table to recognize any new drives added since the initial operating system installation.

IMPORTANT: Although all of the disks connected to the SMW are available to the system, grub detects only the first 16 devices. Therefore, if you add a disk to the SMW after the SMW is connected to the boot RAID, it is advisable to reboot the SMW before continuing this procedure.

- a. Back up the current grub device mapping file.

```
smw:~ # mv /boot/grub/device.map /boot/grub/device.map-YYYYMMDD
```

- b. Invoke the grub utility to create a new device mapping file.

```
smw:~ # grub --device-map=/boot/grub/device.map
Probing devices to guess BIOS drives. This may take a long time.
GNU GRUB version 0.97 (640K lower / 3072K upper memory)
grub> quit
```

The file `/boot/grub/device.map` is now updated to reflect all drives, using the standardized drive naming. To verify the contents of the file:

```
smw:~ # cat /boot/grub/device.map
(fd0) /dev/fd0
(hd0) /dev/sda
(hd1) /dev/sdc
```

6. Create a new file system on the backup drive root partition by executing the mkfs command.

```
smw:~ # mkfs -t ext3 /dev/sdb2
mke2fs 1.41.1 (01-Sep-2008)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
19275776 inodes, 77091918 blocks
3854595 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=4294967296
2353 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
4096000, 7962624, 11239424, 20480000, 23887872, 71663616
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
This filesystem will be automatically checked every 33 mounts or
180 days, whichever comes first. Use tune2fs -c or -i to override.
smw:~ #
```

7. Mount the new backup root file system on /mnt.

```
smw:~ # mount /dev/sdb2 /mnt
```

8. Confirm the running root file system device.

```
smw:~ # df
Filesystem      1K-blocks    Used    Available    Use%    Mounted on
/dev/sda2       303528624  6438700   281671544     3%      /
udev            1030332     116      1030216      1%      /dev
/dev/sdb2       306128812  195568   290505224     1%      /mnt
```

The running root file system device is the one mounted on /.

9. Dump the running root file system to the backup drive.

```
smw:~ # cd /mnt
smw:~ # dump 0f - /dev/sda2 | restore rf -
DUMP: WARNING: no file `/etc/dumpdates'
DUMP: Date of this level 0 dump: Thu Nov 11 06:55:29 2010
DUMP: Dumping /dev/sda2 (/) to standard output
DUMP: Label: none
DUMP: Writing 10 Kilobyte records
DUMP: mapping (Pass I) [regular files]
DUMP: mapping (Pass II) [directories]
DUMP: estimated 4003398 blocks.
DUMP: Volume 1 started with block 1 at: Thu Nov 11 06:57:38 2010
DUMP: dumping (Pass III) [directories]
DUMP: dumping (Pass IV) [regular files]
```

```

restore: ./lost+found: File exists
DUMP: 81.99% done at 10941 kB/s, finished in 0:01
DUMP: Volume 1 completed at: Thu Nov 11 07:04:01 2010
DUMP: Volume 1 4008910 blocks (3914.95MB)
DUMP: Volume 1 took 0:06:23
DUMP: Volume 1 transfer rate: 10467 kB/s
DUMP: 4008910 blocks (3914.95MB)
DUMP: finished in 383 seconds, throughput 10467 kBytes/sec
DUMP: Date of this level 0 dump: Thu Nov 11 06:55:29 2010
DUMP: Date this dump completed: Thu Nov 11 07:04:01 2010
DUMP: Average transfer rate: 10467 kB/s
DUMP: DUMP IS DONE

```

10. To make the backup drive bootable, reinstall the grub boot facility on that drive.

- a. Create a unique file on the backup drive to be used to identify that drive to the grub boot facility.

```

smw:~ # cd /
smw:~ # touch /mnt/THIS_IS_SDx

```

- b. Invoke the grub boot utility. Within grub:

1. Execute the find command to locate the drive designation that grub uses.
2. Select the drive to which the boot blocks will be installed with the root command.
3. Use the setup command to set up and install the grub boot blocks on that drive.

```

smw:~ # grub --no-curses
GNU GRUB version 0.97 (640K lower / 3072K upper memory)
[ Minimal BASH-like line editing is supported. For the first word, TAB^[]
lists possible command completions. Anywhere else TAB lists the possible
completions of a device/filename. ]
grub> find /THIS_IS_SDx
find /THIS_IS_SDx
(hd1,1)
grub> root (hd1,1)
root (hd1,1)
Filesystem type is ext2fs, partition type 0x83
grub> setup (hd1)
setup (hd1)
Checking if "/boot/grub/stage1" exists... yes
Checking if "/boot/grub/stage2" exists... yes
Checking if "/boot/grub/e2fs_stage1_5" exists... yes
Running "embed /boot/grub/e2fs_stage1_5 (hd1)"... 17 sectors are embedded.
succeeded
Running "install /boot/grub/stage1 (hd1) (hd1)1+17 p
(hd1,1)/boot/grub/stage2 /boot/grub/menu.lst"... succeeded
Done.
grub> quit

```

Note that the Linux grub utility and boot system always refer to drives as `hd`, regardless of the actual type of the drives.

11. Unmount the backup root partition.

```

smw:~ # umount /dev/sdb2

```

The drive is now bootable once plugged in or cabled as the primary drive.

Desk-side SMW: Set Up the Bootable Backup Drive as an Alternate Boot Device

The following procedure modifies a bootable backup drive for a desk-side SMW in order to boot from and run the desk-side SMW from the backup root partition.

IMPORTANT: To boot from this backup drive, the primary boot drive must still be operable and able to boot the grub boot blocks installed. If the backup drive is modified to boot as an alternate boot device, it will no longer function as a bootable backup if the primary drive fails.

Desk-side SMW: Set Up the Bootable Backup Drive as an Alternate Boot Device

NOTE: This procedure will not provide a usable backup drive that can be booted in the event of a primary drive failure.



CAUTION: The disk device names shown in this procedure are provided as examples only. Substitute the correct disk devices for your system. For example, on an SMW with three SMW disks, the boot disk is /dev/sda and the bootable backup disk is /dev/sdc. On an SMW with two SMW disks, the boot disk is /dev/sda and the bootable backup disk is /dev/sdb.

1. Mount the backup drive's root partition.

```
smw:~ # mount /dev/sdX2 /mnt
```

2. Create a new boot entry in the /boot/grub/menu.lst file. This entry should be a duplicate of the primary boot entry with the following changes:

- Modify the title to uniquely identify the backup boot entry.
- Modify the `root (hd0,1)` directive to reflect the grub name of the backup drive. If you do not know the grub name of the backup drive, it is provided in the /boot/grub/device.map file on the primary drive.
- Modify the `root=` and `resume=` specifications to reference the backup drive device.

An example /boot/grub/menu.lst file follows. Note the new entry at the end of the file. This example references /dev/sda as the primary drive and /dev/sdc as the backup drive.

```
smw:~ # cat /boot/grub/menu.lst
# Modified by YaST2. Last modification on Wed Dec 9 15:09:52 UTC 2009
default 0
timeout 8
##YaST - generic_mbr
gfxmenu (hd0,1)/boot/message
##YaST - activate

###Don't change this comment - YaST2 identifier: Original name: linux###
title SUSE Linux Enterprise Server 11 - 2.6.27.19-5
    root (hd0,1)
    kernel /boot/vmlinuz-2.6.27.19-5-default root=/dev/sda2 \
    resume=/dev/sda1 splash=silent crashkernel=256M-:128M@16M showopts
vga=0x31a \
    initrd /boot/initrd-2.6.27.19-5-default

###Don't change this comment - YaST2 identifier: Original name: failsafe###
title Failsafe -- SUSE Linux Enterprise Server 11 - 2.6.27.19-5
    root (hd0,1)
    kernel /boot/vmlinuz-2.6.27.19-5-default root=/dev/sda2 showopts \
```

```

ide=nodma apm=off noresume edd=off powersaved=off nohz=off highres=off \
processor.max_cstate=1 x11failsafe vga=0x31a \
initrd /boot/initrd-2.6.27.19-5-default

###Don't change this comment - YaST2 identifier: Original name: floppy###
title Floppy
    rootnoverify (fd0)
    chainloader +1

### New entry allowing a boot of the back-up drive when the primary drive
### is still present.
title BACK-UP DRIVE - SUSE Linux Enterprise Server 11 - 2.6.27.19-5
    root (hd0,1)
    kernel /boot/vmlinuz-2.6.27.19-5-default root=/dev/sdc2 \
    resume=/dev/sdc1 splash=silent crashkernel=256M-:128M@16M showopts
vga=0x31a \
    initrd /boot/initrd-2.6.27.19-5-default

```

3. Modify the backup drive's `/etc/fstab` file to reference the secondary drive slot rather than the first drive slot.
 - a. Examine the backup drive's `fstab` file.

```

smw:~ # cat /mnt/etc/fstab
/dev/sda1 swap swap defaults 0 0
/dev/sda2 / ext3 acl,user_xattr 1 1
proc /proc proc defaults 0 0
sysfs /sys sysfs noauto 0 0
debugfs /sys/kernel/debug debugfs noauto 0 0
usbfs /proc/bus/usb usbfs noauto 0 0
devpts /dev/pts devpts mode=0620,gid=5 0 0

```

- b. Edit the `/mnt/etc/fstab` file, changing `/dev/sda1` and `/dev/sda2` to reference the backup drive. In the following example, the backup drive is `/dev/sdc`.

```

smw:~ # vi /mnt/etc/fstab
/dev/sdc1 swap swap defaults 0 0
/dev/sdc2 / ext3 acl,user_xattr 1 1
proc /proc proc defaults 0 0
sysfs /sys sysfs noauto 0 0
debugfs /sys/kernel/debug debugfs noauto 0 0
usbfs /proc/bus/usb usbfs noauto 0 0
devpts /dev/pts devpts mode=0620,gid=5 0 0

```

4. Unmount the backup drive.

```
smw:~ # umount /dev/sdX2
```

The SMW can now be shut down and rebooted. Upon display of the **Please select boot device** prompt, select the BACK-UP DRIVE - SLES 11 entry to boot the backup root partition.

Configure the Simple Event Correlator (SEC)

The System Management Workstation (SMW) 7.2.UP00 release includes the Open Source simple event correlator (SEC) package, `sec-2.7.0`, and an SEC support package, `cray-sec-4.0.0`. The SEC support

package contains control scripts to manage the starting and stopping of SEC around a Cray mainframe boot session, in addition to other utilities.

If you plan to use the Cray SEC, see *Configuring SEC Software for a Cray XC, Cray XE, or Cray XK System* (S-2542) for configuration procedures.

Configure Boot RAID

Prerequisites and Assumptions for Configuring the Boot RAID

NOTE: Cray ships systems with much of this software installed and configured. You may not need to perform all of the steps described in this chapter unless you are making changes to the configuration.

Cray provides support for system boot RAID from two different vendors: Data Direct Networks, Inc. (DDN) and NetApp, Inc. If you are configuring a DDN boot RAID, follow the procedures in [Configure and Zone the Boot RAID Using a DDN Storage System](#) on page 53. If you are configuring a NetApp, Inc. Engenio™ boot RAID, follow the procedures in [Configure the Boot RAID Using a NetApp, Inc. Engenio Storage System](#) on page 62.

In typical system installations, the RAID provides the storage for both the boot node root file systems and the shared root file system. Although the boot node manages these file systems during normal operation, the SMW performs the initial installation of the CLE operating system and the Cray software packages on the boot RAID disks.

For a system configured for SMW high availability (HA) with the SMW failover feature, the boot RAID is also used for the shared log, MySQL database, and /home file system.

In typical system installations, RAID units provide user and scratch space and can be configured to support a variety of file systems. For more information about configuring RAIDs, see *Managing Lustre for the Cray Linux Environment (CLE)* (S-0010), which is provided with your CLE release package.

Prerequisites for this task are:

- The SMW has an Ethernet connection to the Hardware Supervisory System (HSS) network.
- The boot nodes have Ethernet connections to the SMW.
- The SMW has a Fibre Channel (FC) or Serial Attached SCSI (SAS) connection to the boot RAID or to a FC switch.
- The boot nodes have a FC or SAS connection to the boot RAID or to a FC switch.
- The SDB node has a FC or SAS connection to the boot RAID or to a FC switch.
- If a dedicated `syslog` node is configured, it has a FC or SAS connection to the boot RAID or to a FC switch.
- The DDN RAID uses LUNs (Logical Units). The NetApp, Inc. RAID uses volumes.

The SDB node may have an Ethernet connection. A dedicated `syslog` node does not have an Ethernet connection.

Configure and Zone the Boot RAID Using a DDN Storage System

NOTE: The instructions in this section supersede the documentation supplied by the RAID manufacturer.

You can use either the command-line interface (telnet) or the DDN Silicon Storage Appliance Manager GUI to configure a DDN storage system for your boot RAID.

Configure a DDN Storage System for the Boot RAID via telnet

Configure the boot RAID password when using DDN devices

1. If necessary, log on to the SMW as `crayadm`.
2. From the SMW, use the `telnet` command to log on to the RAID controller.

```
crayadm@smw:~> telnet 10.1.0.1
```

3. Login as `admin` and enter the default password.

```
login: admin
Password:
CAB01-L1:
```

4. Change the password by entering the current password for `admin`, and then enter the information as prompted.

```
CAB01-L1: password
Enter current password:
*****
Enter a new name to replace 'admin', or return to leave unchanged:

Administrative user name 'admin' unchanged.

Enter new password:
*****
Re-enter the new password:
*****

Enter a new name to replace 'user', or return to leave unchanged:

General user name 'user' unchanged.

Enter new password:
*****
Re-enter the new password:
*****

Password for general user 'user' unchanged.

Committing changes.
CAB01-L1:
```

Configure the LUNs for DDN devices

1. On the RAID controller, use the `lun` command to view the existing configuration. If the RAID is already configured, the command returns the current LUN status.

```
CAB01-L1: lun
```

Logical Unit Status							
LUN	Label	Owner	Status	(Mbytes)	Capacity Size	Block Tiers	Block Tier list

2. Use the `lun` command to delete preexisting LUNs.

```
CAB01-L1: lun delete=x where x is the LUN to be deleted.
```

3. Use the `lun` command to add, configure, and format a LUN.

CAB01-L1: `lun add=x` where *x* is the LUN to be configured. This command initiates a dialog similar to the following example.

LUN configuration dialog for DDN devices

```
Enter the LUN (0..127) to add, or 'e' to escape: 0
Enter a label for LUN 0 (up to 12 characters): bootroot0
You can create a single LUN or a LUN group
of smaller LUNs with equal capacity.
Do you want to create a LUN group? (y/N): n
Enter the capacity (in Mbytes) for LUN 0, LUN0
0 for all available capacity (default): 40000
Enter the number of tiers (1..8)
Default will auto select, 'e' to escape: 1
Enter the tiers, each one on a new line, or 'e' to escape: 1
Enter the block size, (512, 1024, 2048, 4096, 8192)
Default is 512, 'e' to escape: 4096
Operation successful: LUN 0 was added to the system.
The LUN must be formatted before it can be used.
Would you like to format the LUN now? (y/N): y
```

The recommended boot RAID LUN configuration is included as a table in *Installing and Configuring Cray Linux Environment (CLE) Software (S-2444)*, which is provided with your CLE release package.

Repeat this step for each LUN that you want to configure and format. The final boot RAID LUN configuration will resemble this example:

```
CAB01-L1: lun
```

Logical Unit Status							
LUN	Label	Owner	Status	(Mbytes)	Capacity Size	Block Tiers	Block Tier list

0	bootroot0	1	Ready	40000	4096	1	1

1	bootroot1	1	Ready	40000	4096	1	2
2	shroot0	1	Ready	280000	4096	1	3
3	sdb	1	Ready	80000	4096	1	2
4	syslog	1	Ready	80000	4096	1	1
5	shroot1	1	Ready	280000	4096	1	4
6	bootroot2	1	Ready	40000	4096	1	3
7	shroot2	1	Ready	280000	4096	1	1

NOTE: On systems that are configured for SMW HA with the SMW failover feature, the boot RAID also includes space for the shared log, MySQL database, and /home file system.

Zone the LUNs for DDN Devices

After the LUNs are configured and formatted, you must grant host access to the LUNs by using a process called *zoning*. Zoning maps a host port on the RAID controller to the LUNs that the host accesses. Although it is possible to allow all hosts to have access to all LUNs, Cray recommends that each host be granted access only to the LUNs it requires.

NOTE: If a LUN is to be shared between failover host pairs, give each host access to the LUN. The SMW host port should be given access to all LUNs.

All zoning relationships are defined with the `zoning` command.

Zoning the LUNs for DDN devices

1. Display the current zoning summary.

```
CAB01-L1: zoning
```

Port Zoning Summary:

Port	World Wide Name	LUN Zoning (External LUN, Internal LUN)
1	21000001FF030759	
2	22000001FF030759	
3	23000001FF030759	
4	24000001FF030759	

If the LUNs are not zoned, this command displays only the Ports and their World Wide Names, the unique identifier in the Fibre Channel storage network.

If the LUNs are zoned, consider clearing some or all of the current settings.

2. Execute the `zoning edit=x` command to begin zoning the LUNs, where *x* is the Port number.

```
CAB01-L1: zoning edit=1
```

This command returns the following prompt:

Enter the new LUN zoning for host port 1.

Enter the unique LUN mapping, as follows:

G.1 GROUP.LUN number
P Place-holder

```

R      Before GROUP.LUN to indicate Read-Only
N      Clear current assignment
<cr>   No change
E      Exit command
?      Display detailed help text

```

```
External LUN 0: is not mapped.  Enter new internal LUN:
```

- The zoning dialog cycles through each LUN in sequence. Enter the LUN number to map the LUN to the selected port, press the Enter key to leave the LUN mapping unchanged, or enter **n** to clear (remove) the current LUN mapping. When you finish mapping LUNs to the selected port, enter **e** to exit. The updated zoning summary displays.

For Port 1, map LUNs 0 through 2 as follows:

```

External LUN 0: is not mapped.  Enter new internal LUN: 0
External LUN 1: is not mapped.  Enter new internal LUN: 1
External LUN 2: is not mapped.  Enter new internal LUN: 2
External LUN 3: is not mapped.  Enter new internal LUN: e
*** Host Port 1: zoning has been updated! ***

```

LUN Zoning				
Port	World Wide Name	(External LUN, Internal LUN)		
1	21000001FF020320	000,000	001,001	002,002
2	22000001FF020320			
3	23000001FF020320			
4	24000001FF020320			

For Port 2, map LUNs 3 through 5, as follows:

```

External LUN 0: is not mapped. Enter new internal LUN: n
External LUN 1: is not mapped. Enter new internal LUN: n
External LUN 2: is not mapped. Enter new internal LUN: n
External LUN 3: is not mapped. Enter new internal LUN: 3
External LUN 4: is not mapped. Enter new internal LUN: 4
External LUN 5: is not mapped. Enter new internal LUN: 5
External LUN 6: is not mapped. Enter new internal LUN: e
*** Host Port 2: zoning has been updated! ***

```

LUN Zoning				
Port	World Wide Name	(External LUN, Internal LUN)		
1	21000001FF020320	000,000	001,001	002,002
2	22000001FF020320	003,003	004,004	005,005
3	23000001FF020320			
4	24000001FF020320			

For Port 3, map LUNs 6 through 8, as follows:

```

External LUN 0: is not mapped. Enter new internal LUN: n
External LUN 1: is not mapped. Enter new internal LUN: n
External LUN 2: is not mapped. Enter new internal LUN: n
External LUN 3: is not mapped. Enter new internal LUN: n
External LUN 4: is not mapped. Enter new internal LUN: n
External LUN 5: is not mapped. Enter new internal LUN: n
External LUN 6: is not mapped. Enter new internal LUN: 6
External LUN 7: is not mapped. Enter new internal LUN: 7
External LUN 8: is not mapped. Enter new internal LUN: 8
External LUN 9: is not mapped. Enter new internal LUN: e
*** Host Port 3: zoning has been updated! ***

```

Port	World Wide Name	LUN Zoning		
		(External LUN, Internal LUN)		
1	21000001FF020320	000,000	001,001	002,002
2	22000001FF020320	003,003	004,004	005,005
3	23000001FF020320	006,006	007,007	008,008
4	24000001FF020320			

4. Enable continuous LUN verification. This process runs in the background with a performance penalty of approximately 1%.

```
CAB01-L1: lun verify=on
[...]
Please enter a LUN ('a' for all LUNs, 'q' to quit): a
All valid LUNs selected
Do you want the verify to run continuously? (y/N): y
```



CAUTION: Turn off LUN verification before you perform maintenance on disk subsystems. If you are running LUN verification and swap or move a back-end channel cable, you can disrupt an entire channel of drives.

5. The final LUN zoning should look like the following example. If it does not, edit any ports using the procedure previously described.

```
CAB01-L1: zoning
```

Port	World Wide Name	LUN Zoning		
		(External LUN, Internal LUN)		
1	21000001FF020320	000,000	001,001	002,002
2	22000001FF020320	003,003	004,004	005,005
3	23000001FF020320	006,006	007,007	008,008
4	24000001FF020320			

6. When you have finished zoning the LUNs, close the `telnet` connection to the RAID controller and return to the SMW.

```
CAB01-L1: logout
```

7. If not already logged on as `root`, `su` to `root`.

```
crayadm@smw:~> su - root
```

8. Reboot the SMW. This enables the SMW to recognize the new LUN configuration and zoning information.

```
smw:~ # reboot
```

After you finish creating, formatting, and zoning the LUNs on the boot RAID, you must partition them. Although partitioning occurs on the SMW, you must use the boot LUN partition table included in *Installing and Configuring Cray Linux Environment (CLE) Software (S-2444)*, which is provided with your CLE release package, because recommendations may change with each CLE release.

Install the Silicon Storage Appliance Manager GUI

Optionally, you can use the Silicon Storage Appliance Manager GUI to configure a DDN storage system for your Boot RAID.

Install the DDN Silicon Storage Appliance Manager

If your Cray system does not have the DDN Silicon Storage Appliance Manager software installed, install it.

The Silicon Storage Appliance Manager installation file (`directGUI_version_Linux_x64_NoVM.bin`) is on a CD that is provided with your system. Before you begin the following procedure, copy the file as `root` user, from the CD to a temporary location on the SMW, for example, `/tmp`.

Identify the installed version of the DDN Silicon Storage Appliance Manager software

You must be logged in as `root` to perform this task.

1. Start the GUI.

```
smw:~ # /usr/local/ddn/Silicon_Storage_Appliance_Manager
```

Or

```
smw:~ # /opt/Silicon_Storage_Appliance_Manager/Silicon_Storage_Appliance_Manager
```

2. Select from Help>About the Silicon Storage Appliance Manager menu bar to identify the version of the GUI. If you are using a Silicon Storage Appliance Manager GUI version earlier than 2.07, download and install the new Silicon Storage Appliance Manager software.

Install the Silicon Storage Appliance Manager software for DDN devices

1. If you are already logged in to the SMW, `su` to `root` and enter the following command:

```
smw:~ # ssh -X
```

Alternatively, if you are not already logged on to the SMW with X windows port forwarding enabled enter the following command:

```
workstation> ssh -X smw
```

2. Move to the directory to which you copied the installation file.
3. Initiate the installation process:

```
smw:~ # ./directGUI_version_Linux_NoVM.bin
```

The Introduction window displays.

4. Select Next.
The License Agreement window displays.
5. Accept the License Agreement and select Next.
The Choose Uninstall Folder window displays.
6. The uninstall folder window displays the installed software. Select the uninstall folder (file path) you want to uninstall and select Next.

- a. Select the complete icon for a complete uninstall.
- b. Select Done.

The Choose Java Virtual Machine window opens.

7. Choose a Java Virtual Machine and select Next.

The Choose Install Folder window opens.

8. Choose the install folder and select Next.

The Choose Link Folder window opens.

9. Choose the link folder and select Next.

The Pre-Installation Summary window opens.

10. Review the pre-installation summary. If everything is correct, select Install.

When the installation completes, the Install Complete window opens.

11. The installation is complete; select Done.

Install the DDN Silicon Storage Appliance Manager

If your Cray system does not have the DDN Silicon Storage Appliance Manager software installed, install it.

The Silicon Storage Appliance Manager installation file (`directGUI_version_Linux_x64_NoVM.bin`) is on a CD that is provided with your system. Before you begin the following procedure, copy the file as `root` user, from the CD to a temporary location on the SMW, for example, `/tmp`.

Install the Silicon Storage Appliance Manager software for DDN devices

1. If you are already logged in to the SMW, `su` to `root` and enter the following command:

```
smw:~ # ssh -X
```

Alternatively, if you are not already logged on to the SMW with X windows port forwarding enabled enter the following command:

```
workstation> ssh -X smw
```

2. Move to the directory to which you copied the installation file.
3. Initiate the installation process:

```
smw:~ # ./directGUI_version_Linux_NoVM.bin
```

The Introduction window displays.

4. Select Next.

The License Agreement window displays.

5. Accept the License Agreement and select Next.

The Choose Uninstall Folder window displays.

6. The uninstall folder window displays the installed software. Select the uninstall folder (file path) you want to uninstall and select Next.
 - a. Select the complete icon for a complete uninstall.
 - b. Select Done.

The Choose Java Virtual Machine window opens.
7. Choose a Java Virtual Machine and select Next.

The Choose Install Folder window opens.
8. Choose the install folder and select Next.

The Choose Link Folder window opens.
9. Choose the link folder and select Next.

The Pre-Installation Summary window opens.
10. Review the pre-installation summary. If everything is correct, select Install.

When the installation completes, the Install Complete window opens.
11. The installation is complete; select Done.

Configure and Zone the LUNs for DDN Devices

You can use the DDN Silicon Storage Appliance Manager to configure and zone LUNs. See the DDN Silicon Storage Appliance Manager GUI documentation for additional information.

Configure Remote Logging of DDN Messages

To avoid overflowing the controller's internal error log, follow this procedure to enable and configure remote logging of DDN messages.

Configure the syslog on DDN storage devices

The message logs for DDN S2A 8500 and S2A 9550 storage devices are written to a fixed area of controller memory. Some storage problems can produce output that exceeds the amount of available memory space in the controller and force initial error messages, which are necessary to isolate the root cause of the problem, to overflow and scroll out of memory.

Cray recommends that you reconfigure the S2A network logging to enable the SMW to store and retrieve failure information.

When you configure RAID controllers to send syslog messages to the SMW, the messages from the RAID are automatically stored in `/var/log/cray/log/raid-yyyymmdd`, assuming the RAID controllers utilize an IP address beginning with `10.1.0`. If your RAID controllers begin with a different IP address, set the `llm_raid_ip` variable in the `SMWinstall.conf` file.

1. Enable syslog on the DDN device to the SMW.

Log on as `admin` to the DDN device. This example uses the command-line interface (`telnet`) and uses the boot RAID and its default IP address.

NOTE: Use the password that you set in [Configure the Boot RAID Password When Using DDN Devices](#).

```
smw:~ # telnet 10.1.0.1
login: admin
Password: *****
```

Verify the current network settings by using the `network` command. The settings for `syslog` appear at the end of the display.

```
CAB01-L1: network
...

Syslog:                                DISABLED
Syslog IP Address:                     0.0.0.0
Syslog Port unit #1:                   514
Syslog Port unit #2:                   514
```

Enable `syslog` to the SMW IP address 10.1.1.1 on the default port of 514.

```
CAB01-L1: network syslog=on
CAB01-L1: network syslogip=10.1.1.1
CAB01-L1: network syslogport=514
```

Verify the new settings.

```
CAB01-L1: network
...

Syslog:                                ENABLED
Syslog IP Address:                     10.1.1.1
Syslog Port unit #1:                   514
```

2. Log off the DDN device.

```
CAB01-L1: logout
```

Configure the Boot RAID Using a NetApp, Inc. Engenio Storage System

NOTE: The instructions in this section apply for both SAS (Serial Attached SCSI) and Fibre Channel RAIDs and supersede the documentation supplied by the RAID manufacturer.

Use the SANtricity Storage Manager utility from NetApp, Inc. to perform the procedures in this section. SANtricity is provided as a separate package and is installed from a CD. It may already be installed on your SMW.

A prerequisite to complete this task is that you are familiar with using the SANtricity interface.

Install SANtricity Storage Manager Software for NetApp, Inc. Engenio Devices

The SANtricity Storage Manager software is generally preinstalled and the SANtricity media is shipped with the system. However, if your Cray system does not have the software installed on the SMW, you can install it. The SANtricity SMClient executable is found in `/opt/SMgr/client`.

Install the SANtricity Software

1. If you are not already logged on as `root`, `su` to `root`.

```
crayadm@smw:~> su - root
```

2. If you are installing from the SANtricity Storage Manager CD, insert it into the SMW CD drive. Verify that the media has mounted automatically; if not, mount it manually.

```
smw:~ # mount /dev/cdrom /media/cdrom
```

Or, if you are installing from the `SMIA-LINUX-10.70.A0.25.bin` file, copy `SMIA-LINUX-10.70.A0.25.bin` to `/home/crayadm`.

```
smw:~ # cp ./SMIA-LINUX-10.70.A0.25.bin /home/crayadm/
```

3. Set the `DISPLAY` environment variable.

```
smw:~ # export DISPLAY=:0.0
```

4. Verify that the X Window System is functioning by launching `xterm` or executing the `xlogo` utility.

```
smw:~ # xterm
```

or:

```
smw:~ # xlogo
```

Then, exit the Xlogo window or `xterm`.

5. Invoke the executable file.

If you are installing from the CD:

```
smw:~ # /bin/bash /media/cdrom/install/SMIA-LINUX-10.70.A0.25.bin
```

Or, if you are installing from a directory:

```
smw:~ # /home/crayadm/SMIA-LINUX-10.70.A0.25.bin
```

6. Select **Next**. The **License Agreement** window displays.
7. Accept the license agreement and select **Next**. The **Select Installation Type** window displays.
8. Select **Typical (Full Installation)**, then select **Next**.
The **Multi-Pathing Driver Warning** window displays.
9. Select **OK**. The **Pre-Installation Summary** window displays.

10. Select **Install**.

The **Installing SANtricity** window displays and shows the installation progress. When the installation completes, an **Install Complete** window appears.

11. Select **Done**. The SANtricity client is installed in `/usr/bin/SMclient` and is currently running.

12. To execute `SMclient` from the `crayadm` account, change the ownership and permissions for executable files; otherwise, you will need to execute `SMclient` as `root`.

```
smw:~ # cd /opt
smw:~ # chown crayadm SMgr
smw:~ # chmod 775 SMgr
smw:~ # cd SMgr/client
smw:~ # chmod 755 SMcli SMclient
smw:~ # cd /var/opt
smw:~ # chown -R crayadm:crayadm SM
smw:~ # chmod -R ug+w SM
```

13. Close the file browser and eject the CD.

```
smw:~ # eject
```

Configure the LUNs for NetApp, Inc. Engenio Devices

Create the Volume Group and the LUNs that are members of it.

Create the boot RAID volume group for NetApp, Inc. Engenio devices

Create the 3+1 Volume Group and 1 Global Hot Spare across the first five disks for a 1.5 TB Volume Group. The **Array Management** window should still be displayed after performing the procedure.

You must be logged on to the SMW as `crayadm` to perform this task.

1. Start the SANtricity Storage Manager.

```
crayadm@smw:~> /usr/bin/SMclient
```

The SANtricity Storage Manager window appears.

2. If the **Select Addition Method** window appears, choose one of the following options; otherwise, continue with the next step.

- **Automatic** - Select this option if you did not assign IP addresses to the storage array controllers using a serial connection. The SANtricity software automatically detects the available controllers, in-band, using the Fibre Channel link.
- **Manual** - Select this option if you have already assigned IP addresses to the storage array controllers.

NOTE: The following steps assume that you selected the **Manual** option.

3. Double-click the name for the storage array that you want to configure. The **Array Management** window displays.

4. Select the **Logical/Physical** tab.

5. Right-click **Unconfigured Capacity** and select **Create Volume**. The **Create Volume** wizard displays.
 - a. Select **Next** on the **Introduction (Create Volume)** window.
 - b. Select the **Manual** option on the **Specify Volume Group (Create Volume)** window.
 - c. Select tray 85, slots 1-4 and select **Add**.
 - d. Verify that the RAID level is set to 5.
 - e. Select **Calculate Capacity**.
 - f. Select **Next** on the **Specify Volume Group (Create Volume)** window.

When you create the first **Volume Group**, you are prompted to create the first volume.

Create and Configure Volumes for NetApp, Inc. Engenio Devices

Configure the boot RAID with enough LUNs to support the various system management file systems. (Cray recommends a minimum of nine LUNs.)

Use the boot LUN configuration table included in *Installing and Configuring Cray Linux Environment (CLE) Software (S-2444)*, which is provided with your CLE release package, as a guideline for configuration information for each volume.

NOTE: For a system configured for SMW HA with the SMW failover feature, the boot RAID must also include space for the shared log, MySQL database, and /home file system.

1. Enter a new volume capacity. Specify units as GB or MB.
2. Enter a name for the volume.
3. Select the **Customize Settings** option.
4. Select **Next** in the **Specify Capacity/Name (Create Volume)** window.
5. Verify the settings on the **Customize Advanced Volume Parameters (Create Volume)** window. These settings are used for the all of the LUNs.
 - For **Volume I/O characteristics type**, verify that **File System** is selected.
 - For **Preferred Controller Ownership**, verify that **Slot A** is selected. This places the LUN on the A Controller.
6. Select **Next** in the **Customize Advanced Volume Parameters (Create Volume)** window.
7. In the **Specify Volume to LUN Mapping** window, select the **Default** mapping option.
8. For **Host** type, select **Linux** from the drop-down menu.
9. Select **Finish** in the **Specify Volume to LUN Mapping** window.
10. When prompted to create more LUNs in the **Creation Successful (Create Volume)** window, select **Yes** unless this is the last volume you are creating. If this is the last volume, select **No** and skip to [14](#) on page 66

11. In the **Allocate Capacity (Create Volume)** window, verify that **Free Capacity** is selected on **Volume Group 1 (RAID 5)**.
12. Select **Next** in the **Allocate Capacity (Create Volume)** window.
13. Repeat steps 1 through 13 to create all of the volumes described in the boot LUN configuration table in *Installing and Configuring Cray Linux Environment (CLE) Software (S-2444)*, which is provided with your CLE release package.
14. Select **OK** in the **Completed (Create Volume)** window.
15. Create a hot spare. The hot spare provides a ready backup if any of the drives in the Volume Group fail.
 - a. Right-click on the last drive in the slot 14 icon on the right portion of the window and select **Hot Spare Coverage**.
 - b. Select the **Manually Assign Individual Drives** option.
 - c. Select **OK**.
 - d. Select **Close**.
16. Exit the tool.

Configure Remote Logging of NetApp, Inc. Engenio Storage System Boot RAID Messages

The NetApp, Inc. Engenio storage system uses SNMP to provide boot RAID messages. See your NetApp, Inc. Engenio Storage System documentation for additional information.

Zone the QLogic™ FC Switch

If you have a QLogic Fibre Channel Switch, follow [Configure zoning for a QLogic SANbox switch using QuickTools utility](#) on page 67 to zone the LUNs on your QLogic SANBox™ switch by using a utility called *QuickTools*.

NOTE: If a LUN is to be shared between failover host pairs, each host must be given access to the LUN. The SMW host port should be given access to all LUNs.

QuickTools is an application that is embedded in the QLogic switch and is accessible from a workstation browser with a compatible Java™ plug-in. You must have a Java browser plugin, version 1.4.2 or later.

These instructions assume that the disk device has four host ports connected to ports 0-3 for the QLogic SANbox switch. The following connections are also required:

- The SMW must be connected to port 10 on the SANBox.
- The boot node must be connected to port 4 on the SANBox.
- The SDB node must be connected to port 5 on the SANBox.
- If a dedicated syslog node is configured, it must be connected to port 6 on the SANBox.

Zoning is implemented by creating a *zone set*, adding one or more zones to the zone set, and selecting the ports to use in the zone.

Prerequisite for this task: the SANBox is configured and on the HSS network.

Configure zoning for a QLogic SANbox switch using QuickTools utility

1. Start a web browser.
2. Enter the IP address of your switch. If your configuration has a single switch, the IP address is 10.1.0.250. The IP address of each RAID controller is preconfigured by Cray and is listed on a sticker on the back of the RAID controller.
3. Enter the login name and password when the Add a New Fabric window pops up and prompts for them. The default administrative login name is `admin`, and the default password is `password`.
4. The QuickTools utility displays in your browser. Select Add Fabric.
NOTE: If you receive a dialog box notification that the request failed to connect over a secured connection, select Yes and continue.
5. The switch is located and displayed in the window. Double-click the switch icon. Information about the switch displays in the right panel.
6. At the bottom of the panel, select the Configured Zonesets tab.
7. From the toolbar menu, select Zoning and then Edit Zoning. The Edit Zoning window displays.
8. Select the Zone Set button. The Create a Zone Set window displays. Create a new zone set. (In this example, assume that the zone set is named `XT0`.)
9. Right-click the XT0 zone and select Create a Zone.
10. Create a new zone named BOOT.
11. On the right panel, select the button in front of BOOT to open a view of the domain members.
12. Ports 0, 4, 5, and 10 are added to the BOOT zone. Define the ports in the zone to ensure that the discovery of LUNs is consistent among the SMW, the boot node, and the SDB node.
 - a. Using the mouse, left-click Port # 0 and drag it to the BOOT zone.
 - b. Using the mouse, left-click Port # 4 and drag it to the BOOT zone. This port is for the boot node.
 - c. Using the mouse, left-click Port # 5 and drag it to the BOOT zone. This port is for the SDB node.
 - d. Using the mouse, left-click Port # 10 and drag it to the BOOT zone. This port is for the SMW.
13. Select **Apply**. The error-checking window displays.
14. When prompted, select **Perform Error Check**.
15. After confirming that no errors were found, select **Save Zoning**.

16. When prompted to activate a Zone Set, select **Yes** and then select the appropriate XT0 zone set.

At this point, Cray recommends that you create a backup of your switch configuration ([Create a backup of your QLogic switch configuration](#) on page 68) before you close and exit the application.

Create a backup of your QLogic switch configuration

Create a backup of your QLogic switch configuration with the QuickTools utility. You must have a Java browser plugin, version 1.4.2 or later to use QuickTools.

If you need to start your web browser and open the QuickTools utility, complete steps 1 through 4. If you currently have the QuickTools utility open, skip to step 5.

1. Start a web browser.
2. Enter the IP address of your switch. The IP address of each RAID controller is preconfigured by Cray and is listed on a sticker on the back of the RAID controller.
3. Enter the login name and password when the **Add a New Fabric** window pops up and prompts for them. The RAID default administrative login name is `admin`, and the default password is `password`.
4. The QuickTools utility appears. Select **Add Fabric**. If you receive a dialog box that states that the request failed to connect over a secured connection, select **Yes** and continue.
5. From within the QuickTools utility, complete the configuration backup:
 - a. At the top bar, select **Switch** and then **Archive**. A **Save** window pops up with blanks for **Save in:** and **File Name:**.
 - b. Enter the directory (for example, `crayadm`) and a file name (for example, `sanbox_archive`) for saving your QLogic switch configuration.
 - c. Select the **Save** button.
6. Close and exit the application.

Rediscover the LUNs

This procedure causes the SMW to rediscover the LUNs and zones that you created.

Reboot the SMW and verify LUNs are recognized

1. Log on as the `root` user.

```
crayadm@smw:~> su - root
```

2. Enter the following command to ensure that the LUNs are recognized:

```
smw:~ # reboot
```



CAUTION: Failure to reboot the SMW at this point could produce unexpected results later on.

3. Log on as the `root` user.

```
crayadm@smw:~> su - root
```

4. Execute the `lsscsi` command to verify that the LUNs (volumes) have been rediscovered.

```
smw:~ # lsscsi
```

5. List the disk devices by using the `fdisk` command to verify that the LUNs (volumes) are configured according to the boot LUN configuration table in *Installing and Configuring Cray Linux Environment (CLE) Software* (S-2444), which is provided with your CLE release package.

```
smw:~ # fdisk -l
```

Partition the LUNs

After you finish creating, formatting, and zoning the LUNs on the boot RAID, you must partition them. For procedures and LUN partitioning recommendations, see *Installing and Configuring Cray Linux Environment (CLE) Software* (S-2444), which is provided with your CLE release package.

Set Up Boot RAID Space for Direct-attached Lustre

If your system will use Direct-attached Lustre (DAL) you need to create LUNs for DAL, as well as a LUN and file system in the Image Management and Provisioning System (IMPS) directory, `/var/opt/cray/imps`. This must be done before installing CLE and DAL.

Create LUNs and an IMPS file system for DAL

1. Follow the instructions in [Prerequisites and Assumptions for Configuring the Boot RAID](#) on page 53 and [Configure and Zone the Boot RAID Using a DDN Storage System](#) on page 53 to create and configure LUNs on your specific Boot RAID storage device. The LUN for IMPS must be at least 50 GB; Cray recommends 250 GB.
2. Log on to the SMW as `root`.

```
crayadm@smw:~> su - root
```

3. Use the `lsscsi` command to verify that the LUN was recognized. The output shows the ATA SMW devices, followed by the non-SMW device or devices. This example output is for an SMW with four internal SATA drives and the LUN is `/dev/sde`. The steps that follow use this example LUN.

```
smw:~ # lsscsi
[0:0:0:0]    disk      ATA      FUJITSU MHZ2160B 8A22  /dev/sda
```

```
[0:0:1:0]    disk    ATA      ST91000640NS    AA02    /dev/sdb
[0:0:2:0]    disk    ATA      FUJITSU MHZ2160B 8A22    /dev/sdc
[0:0:3:0]    disk    ATA      FUJITSU MHZ2160B 8A22    /dev/sdd
[1:0:0:0]    disk    LSI      INF-01-00      0786    /dev/sde
```

4. Use the `fdisk` command to create a partition for IMPS that spans the entire LUN. Type `w` to write the changes and exit `fdisk`.

```
smw:~ # fdisk /dev/sde
smw:~ # w
```

5. Use the `mkfs` command to create an `ext3` file system in the partition.

```
smw:~ # mkfs.ext3 /dev/sde1
```

6. Create a mount point.

```
smw:~ # mkdir -p /var/opt/cray/imps
```

7. Add the file system entry to `/etc/fstab`.

- a. Determine the `by-id` persistent device name.

```
smw:~ # ls -l /dev/disk/by-id/* | grep sde1
lrwxrwxrwx 1 root root 10 Jan 27 16:20 /dev/disk/by-id/
scsi-360080e500023d45e000006ab510a\
5817-part1 -> ../../sde1
```

- b. Edit `/etc/fstab` to add the device path.

```
/dev/disk/by-id/scsi-360080e500023d45e000006ab510a5817-part1 /var/opt/cray/
imps/ \
ext3 defaults 0 0
```

8. Mount the IMPS directory.

```
smw:~ # mount /var/opt/cray/imps
```

9. Exit from `root`.

```
smw:~ # exit
crayadm@smw:~>
```

Update or Upgrade Your Cray SMW Software

Cray provides periodic updates to each System Management Workstation (SMW) release, as well as upgrade releases. In an upgrade release, the major and/or minor version number will change, for example from 7.1.UP01 to 7.2.UP00. In an update release only the numbers following the letters UP will have changed.

Follow the procedures in this chapter to install an `SMW 7.2UPxx` package. The procedures provided in this chapter do not change the base operating system version running on your SMW.



CAUTION: You must be running the SUSE Linux Enterprise Server version 11 Service Pack 3 (SLES 11 SP3) SMW base operating system and a release of SMW 7.1 or later on your SMW in order to perform the procedures in this chapter.

NOTE: In the following sections, some examples are left-aligned to better fit the page. Left-alignment has no special significance.

Before You Start

- Determine which SLES version is running on your SMW by executing the following command:

```
crayadm@smw:~> cat /etc/SuSE-release
```

- Read the *SMW README* and *SMW Errata* that are included in the SMW update directory, to determine whether there are any changes to the procedures in this chapter.
- Read the Field Notices (FNs) related to kernel security fixes and apply any needed changes before continuing with the installation.
- For a rack-mount SMW: Enable an Integrated Dell Remote Access Controller 6 (iDRAC6). See [Enable an Integrated Dell Remote Access Controller \(iDRAC6\) on a rack-mount SMW](#) on page 103.
- If you have made local changes to the file `/opt/cray/hss/default/etc/sedc_srv.ini`, note the following information.
- Cray software manages this file as a symbolic link to `/opt/cray/hss/default/etc/sedc_srv.ini.xtek` (Cray XE and Cray XK systems only) or to `/opt/cray/hss/default/etc/sedc_srv.ini.cascade` (Cray XC30 systems only). The following actions are taken during software updates:
 - If the symbolic link exists, it is not altered.
 - If the symbolic link does not exist, it is created as specified above.
 - If `sedc_srv.ini` exists not as a symbolic link but as a regular file, it is renamed to `/opt/cray/hss/default/etc/sedc_srv.ini-YYYYMMDDhhmmss` (where `YYYYMMDDhhmmss` is the date and time the file was renamed) and a new symbolic link is created.

- Before beginning the upgrade, copy `sedc_srv.ini` to a new site-specific file and change the symbolic link to point to that file. After the software update, compare the local file to the distribution `sedc_srv.ini.xtek` or `sedc_srv.ini.cascade` file for any possible changes that should be merged into the locally modified file.
- If you have made local changes to any automation files, such as `/opt/cray/hss/default/etc/auto.xtshutdown`, back them up before beginning the SMW upgrade.
- For Cray XC Systems: If you created any power management profiles that you wish to retain, back up all of the files in the `/opt/cray/hss/7.1.0/pm/profiles` directory before beginning the SMW upgrade.
- For liquid-cooled Cray XC30 Systems only: One or more patches to the SMW 7.1.UP01 release may have installed an `hss.ini` file on your system. This file, and a file containing the default values, `hss.ini.dist` are located in `/opt/tftpboot/ccrd`. If you have not made any site-specific changes to the `hss.ini`, delete this file before beginning the update installation and an active `hss.ini` file will be created as part of the installation. If you have made site-specific changes to this file, save your local copy of `hss.ini` to another location. After installation is complete, copy `hss.ini.dist` to `hss.ini` and re-create your local settings in the new `hss.ini` file.
- If you use the Cray simple event correlator (SEC) and have made local changes to the configuration file `/opt/cray/default/SEC_VARIABLES`, make a backup copy of this file before you begin the upgrade or update. For more information, see *Configuring SEC Software for a Cray XC, Cray XE, or Cray XK System* (S-2542).

NOTE: If you have locally modified versions of `/etc/syslog-ng/syslog-ng.conf` or `/etc/rsyslog.conf`, the local modifications are saved during the upgrade procedure. Also, if you have a locally modified version of `/opt/cray/hss/default/etc/xtnlrd.ini` and the new release includes an updated `xtnlrd.ini` file, your version of the file is preserved during the upgrade and the new file is installed as `xtnlrd.ini.rpmnew`. After the upgrade, compare the two files and merge any changes into your local `xtnlrd.ini` file.

Back Up Your Current Software

Before you install the update package, back up your current SMW software.

Rack-mount SMW: Create an SMW bootable backup drive

IMPORTANT: The disk device names shown in this procedure are only examples. You should substitute the actual disk device names for your system. The boot disk is `phy7` and is slot 0, and the bootable backup disk is `phy6` and is slot 1.



CAUTION: Shut down the Cray system before you begin this procedure.

Also be aware that there may be a considerable load on the SMW while creating the SMW bootable backup drive.

1. Log on to the SMW as `crayadm` and `su` to `root`.

```
crayadm@smw:~> su -
Password:
smw:~ # ls -al /dev/disk/by-path
total 0
drwxr-xr-x 2 root root 380 Mar 15 13:21 .
```

```

drwxr-xr-x 6 root root 120 Mar 11 18:42 ..
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:00:11.0-scsi-0:0:0:0 -> ../../sr0
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:00:12.2-usb-0:3:1.0-scsi-0:0:0:0 -
> ../../sdf
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:00:12.2-usb-0:3:1.1-scsi-0:0:0:0 -
> ../../srl
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:00:12.2-usb-0:3:1.1-scsi-0:0:0:1 -
> ../../sdg
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:00:13.2-usb-0:2:1:1.0-scsi-0:0:0:0 -
> ../../sde
lrwxrwxrwx 1 root root 10 Mar 11 18:42 pci-0000:00:13.2-usb-0:2:1:1.0-scsi-0:0:0:0-
part1 -> ../../sde1
lrwxrwxrwx 1 root root 10 Mar 11 18:42 pci-0000:00:13.2-usb-0:2:1:1.0-scsi-0:0:0:0-
part2 -> ../../sde2
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:05:00.0-sas-phy4-...
lrwxrwxrwx 1 root root 10 Mar 14 15:57 pci-0000:05:00.0-sas-phy4-...
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:05:00.0-sas-phy5-...
lrwxrwxrwx 1 root root 10 Mar 14 16:00 pci-0000:05:00.0-sas-phy5-...
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:05:00.0-sas-phy6-...
lrwxrwxrwx 1 root root 10 Mar 15 13:21 pci-0000:05:00.0-sas-phy6-...
lrwxrwxrwx 1 root root 10 Mar 15 13:21 pci-0000:05:00.0-sas-phy6-...
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:05:00.0-sas-phy7-...
lrwxrwxrwx 1 root root 10 Mar 11 18:42 pci-0000:05:00.0-sas-phy7-...
lrwxrwxrwx 1 root root 10 Mar 11 18:42 pci-0000:05:00.0-sas-phy7-...

```

2. Standardize the SMW's boot-time drive names with the Linux run-time drive names.

IMPORTANT: If the SMW configuration files on the SMW root drive have been modified already (because your site has completed this step at least once after installing your updated SMW base operating system), skip to [3](#) on page 76 ; otherwise, complete this step to standardize the SMW's boot-time drive names with the Linux run-time drive names.

Set up ordered drives on your rack-mount SMW.

- Identify the installed SMW drive model numbers, [3](#) on page 76 serial numbers, and associated Linux device (/dev) names.

Execute the `smwmapdrives` script on the SMW to identify local (internal) drives mounted in the SMW and provide their Linux device (/dev) names.

NOTE: Effective with the SMW 7.2.UP00 release, the `smwmapdrives` script was provided both as a separate file in the release and in the base operating system RPM. If you are running that release or a later one, you can use the installed version of the script to back up the SMW.

```

smw:~ # ./smwmapdrives
List of SMW-installed disk drives
-----
Physical slot 0:
    /dev/sda
    /dev/disk/by-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RDS
    /dev/disk/by-id/scsi-SATA_FUJITSU_MHZ2160_K85DTB227RDS
    /dev/disk/by-path/pci-0000:05:00.0-sas-phy7-0x4433221107000000-lun-0
Physical slot 1:
    /dev/sdc
    /dev/disk/by-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RD7
    /dev/disk/by-id/scsi-SATA_FUJITSU_MHZ2160_K85DTB227RD7
    /dev/disk/by-path/pci-0000:05:00.0-sas-phy6-0x4433221106000000-lun-0
Physical slot 2:
    /dev/sdd
    /dev/disk/by-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RF3
    /dev/disk/by-id/scsi-SATA_FUJITSU_MHZ2160_K85DTB227RF3
    /dev/disk/by-path/pci-0000:05:00.0-sas-phy5-0x4433221105000000-lun-0
Physical slot 3:
    /dev/sdb
    /dev/disk/by-id/ata-ST9500620NS_9XF0665V

```

```

/dev/disk/by-id/scsi-SATA_ST9500620NS_9XF0665V
/dev/disk/by-path/pci-0000:05:00.0-sas-phy4-0x4433221104000000-lun-0
Physical slot 4:
    NOT INSTALLED
Physical slot 5:
    NOT INSTALLED

```

The device names for `by-id` are persistent and will reference the drive, regardless of the slot in which the drive is installed.

`by-path` names reference a physical drive slot only and do not identify the drive installed in that slot. This is the naming used by default for the logging and database drives when the SMW was installed. This `by-path` name is used to specifically install logging and database file systems because the `by-id` device names refer to the physical drive slots expected to be used for those file systems and are provided as the default examples in the SMW installation configuration process.

The `/dev/sdX` drive names are not persistent, except for the `/dev/sda` drive name, which always refers to the drive installed in physical slot 0 (zero) of the SMW. The other `/dev/sdX` names can change with each SMW boot and will change if drives are added, removed, or reordered in the SMW slots. For this reason, the `/dev/sda` drive name can only be used for the rack-mount SMW.

Choose either the `by-id` naming or the `by-path` naming as the site administrative policy for managing the SMW-install disk drives. The following documentation provides the steps necessary to implement this selection on the SMW prior to creating an SMW bootable backup drive.

- b. Back up the following files before proceeding:

```

smw# cp -p /boot/grub/device.map /boot/grub/device.map-YYYYMMDD
smw# cp -p /boot/grub/menu.lst /boot/grub/menu.lst-YYYYMMDD
smw# cp -p /etc/fstab /etc/fstab-YYYYMMDD

```

- c. Edit the grub `device.map` file to reflect physical drive locations.

To provide a direct mapping of the SMW disk drive physical slots to the boot loader (BIOS and grub) drive names, the `device.map` mapping file used by grub should be replaced. Perform the following steps to install new `device.map` file entries to effect this mapping.

1. Edit the grub `device.map` file.
2. Delete all lines.
3. Enter the following lines into the file. These lines show each drive slot's physical location mapped to its boot-time `hd?` name.

NOTE: `by-id` names should not be used in the `device.map` file.

```

# Dell Rackmount r805 SMW
# grub(8) device mapping for boot-drive identification
# hd? numbers are being mapped to their physical
(hd0) /dev/disk/by-path/pci-0000:05:00.0-sas-
phy7-0x4433221107000000-lun-0
(hd1) /dev/disk/by-path/pci-0000:05:00.0-sas-
phy6-0x4433221106000000-lun-0
(hd2) /dev/disk/by-path/pci-0000:05:00.0-sas-
phy5-0x4433221105000000-lun-0
(hd3) /dev/disk/by-path/pci-0000:05:00.0-sas-
phy4-0x4433221104000000-lun-0
(hd4) /dev/disk/by-path/pci-0000:05:00.0-sas-
phy3-0x4433221103000000-lun-0
(hd5) /dev/disk/by-path/pci-0000:05:00.0-sas-
phy2-0x4433221102000000-lun-0

```

- d. Modify the SMW boot drive `/etc/fstab` file to use `by-id` or `by-path` naming.

Modify the SMW file system mounting configuration file to use SMW disk `by-id` or `by-path` naming. Complete this step to replace any `/dev/sdX` disk partition references.

NOTE: Use the output of the `smwmapdrives` script in as a reference for drive names.

Edit `/etc/fstab`, replacing drive `/dev/sdX` references with either the `by-id` or `by-path` name's corresponding device name.

When a reference to `/dev/sda1` is being replaced, replace it with the corresponding "partition" file system suffixed with `-part1`. File system partitions for `/dev/sda` are indicated by the numeral appended to the device name; for example, `/dev/sda1` refers to partition 1 on `/dev/sda`. Replace it with the `by-id` and `by-path` device names and suffix the device name with `-part1`, keeping the same numeral identification.

For example, if the root and swap file systems are currently configured to mount `/dev/sda2`, they should be changed. Using the `by-id` device name from the example in , the `fstab` lines would change from:

```
/dev/sda1 swap                swap                defaults            0 0
/dev/sda2 /                    ext3                acl,user_xattr      1 1to:
/dev/disk/by-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RDS-part1 swap
swap defaults                0 0
/dev/disk/by-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RDS-part2 /
ext3 acl,user_xattr 1 1
```

- e. Modify `/boot/grub/menu.lst` to reflect the device.map BIOS/boot-up drive changes for the `sdX` remapping.

The same device name replacement performed on `/etc/fstab` should also be performed on the grub bootloader `/boot/grub/menu.lst` configuration file. All references to `/dev/sdX` devices should be replaced with corresponding `by-id` or `by-path` device names.

- f. Invoke the grub utility to reinstall the SMW boot loader on the primary boot drive.

Once the changes to `device.map`, `fstab`, and `menu.lst` have been completed, the grub bootloader boot blocks must be updated to reflect changes to the device names. Complete this step to update the boot loader on the boot drive.

Invoke the grub utility and reinstall SMW root-drive boot blocks.

```
smw:~ # grub --no-curses
GNU GRUB version 0.97 (640K lower / 3072K upper memory)
[ Minimal BASH-like line editing is supported. For the first word, TAB
  lists possible command completions. Anywhere else TAB lists the possible
  completions of a device/filename. ]
grub> root (hd0,1)
      root (hd0,1)
      Filesystem type is ext2fs, partition type 0x83
grub> setup (hd0)
      Checking if "/boot/grub/stage1" exists... yes
      Checking if "/boot/grub/stage2" exists... yes
      Checking if "/boot/grub/e2fs_stage1_5" exists... yes
      Running "embed /boot/grub/e2fs_stage1_5 (hd0)"... 17 sectors \
are embedded. Succeeded
      Running "install /boot/grub/stage1 (hd0) (hd0)1+17 p (hd0,1)/boot/grub/
stage2 \
/boot/grub/menu.lst"... succeeded
```

```
Done.
grub> quit
```

3. If the backup drive disk partition table already exists and the partition table on the backup drive matches the partition table that is on the primary boot drive, skip this step; otherwise, create the backup drive disk partition table.

In this example, the partition table consists of two slices. Slice 1 is a 4 GB Linux swap partition. Slice 2 is the balance of disk space used for the root file system.

- a. Use the `fdisk` command to display the boot disk partition layout.

```
smw:~ # fdisk -lu \
/dev/disk/by-path/pci-0000:05:00.0-sas-phy7-0x4433221107000000-lun-0
Disk /dev/disk/by-path/pci-0000:05:00.0-sas-phy7-0x4433221107000000-lun-0: 250.0 GB, \
268435456000 bytes
255 heads, 63 sectors/track, 19457 cylinders, total 312581808 sectors
Units = sectors of 1 * 512 = 512 bytes
Disk identifier: 0x00000082

    Device
/dev/disk/by-path/pci-0000:05:00.0-sas-phy7-0x4433221107000000-lun-0-part1 \
    Boot    Start        End      Blocks   Id  System
           63    16771859    8385898+   82    Linux swap / Solaris
/dev/disk/by-path/pci-0000:05:00.0-sas-phy7-0x4433221107000000-lun-0-part2 \
    Boot    Start        End      Blocks   Id  System
    * 16771860 312576704 147902422+   83    Linux
```

- b. Use the `fdisk` command to configure the bootable backup disk partition layout. Set the bootable backup disk partition layout to match the boot disk partition layout. First, clear all of the old partitions using the `d` command within `fdisk`; next create a Linux swap and a Linux partition; and then write your changes to the disk. For help, type `m` within `fdisk`.

```
smw:~ # fdisk -u \
/dev/disk/by-path/pci-0000:05:00.0-sas-phy6-0x4433221106000000-lun-0

The number of cylinders for this disk is set to 19457.
There is nothing wrong with that, but this is larger than 1024,
and could in certain setups cause problems with:
 1) software that runs at boot time (e.g., old versions of LILO)
 2) booting and partitioning software from other OSs
   (e.g., DOS FDISK, OS/2 FDISK)

Command (m for help): p

Disk /dev/disk/by-path/pci-0000:05:00.0-sas-phy6-0x4433221106000000-lun-0: 250.0 GB, \
268435456000 bytes
255 heads, 63 sectors/track, 19457 cylinders, total 312581808 sectors
Units = sectors of 1 * 512 = 512 bytes
Disk identifier: 0x00000080

    Device
/dev/disk/by-path/pci-0000:05:00.0-sas-phy6-0x4433221106000000-lun-0-part1 \
    Boot    Start        End      Blocks   Id  System
           63    16771859    83828    82    Linux
swap / Solaris
Partition 1 does not end on cylinder boundary.
/dev/disk/by-path/pci-0000:05:00.0-sas-phy6-0x4433221106000000-lun-0-part2 \
    Boot    Start        End      Blocks   Id  System
    167719 312581807 156207044+   83    Linux

Command (m for help): d
Partition number (1-4): 2
```

```

Command (m for help): d
Selected partition 1

Command (m for help): n
Command action
  e   extended
  p   primary partition (1-4)
p
Partition number (1-4): 1
First sector (63-312581807, default 63): (Press the Enter key)
Using default value 63
Last sector, +sectors or +size{K,M,G} (63-312581807, default 312581807): 16771859
Command (m for help): t
Selected partition 1
Hex code (type L to list codes): 82
Changed system type of partition 1 to 82 (Linux swap / Solaris)

Command (m for help): n
Command action
  e   extended
  p   primary partition (1-4)
p
Partition number (1-4): 2
First sector (16771860-312581807, default 16771860): (Press the Enter key)
Using default value 16771860
Last sector, +sectors or +size{K,M,G} (16771860-312581807, default 312581807): (Press
the Enter key)
Using default value 312581807

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.

```

- c. Display the boot backup disk partition layout and confirm it matches your `phy7` sector information.

```

smw:~ # fdisk -lu \
/dev/disk/by-path/pci-0000:05:00.0-sas-phy6-0x4433221106000000-lun-0
Disk /dev/disk/by-path/pci-0000:05:00.0-sas-phy6:1-0x4433221106000000:0-lun0: 250.0
GB, \
268435456000 bytes
255 heads, 63 sectors/track, 19457 cylinders, total 312581808 sectors

```

4. Initialize the swap device.

```

smw:~ # mkswap \
/dev/disk/by-path/pci-0000:05:00.0-sas-phy6-0x4433221106000000-lun-0-part1
mkswap: /dev/disk/by-path/pci-0000:05:00.0-sas-phy6:1-0x4433221106000000:0-lun0-part1:
warning: don't erase bootbits sectors
        (DOS partition table detected). Use -f to force.
Setting up swapspace version 1, size = 8385892 KiB
no label, UUID=c0ef22ac-b405-4236-855b-e4a09b6e94ed

```

5. Create a new file system on the backup drive root partition by executing the `mkfs` command.

```

smw:~ # mkfs -t ext3 \
/dev/disk/by-path/pci-0000:05:00.0-sas-phy6-0x4433221106000000-lun-0-part2
mke2fs 1.41.9 (22-Aug-2009)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
9248768 inodes, 36976243 blocks
1848812 blocks (5.00%) reserved for the super user
First data block=0

```

```

Maximum filesystem blocks=4294967296
1129 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000, 7962624, 11239424, 20480000, 23887872

Writing inode tables:   done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 37 mounts or
180 days, whichever comes first.  Use tune2fs -c or -i to override.

```

6. Mount the new backup root file system on /mnt.

```

smw:~ # mount \
/dev/disk/by-path/pci-0000:05:00.0-sas-phy6-0x4433221106000000-lun-0-part2 /mnt

```

7. Confirm that the backup root file system is mounted.

```

smw:~ # df
Filesystem            1K-blocks      Used Available Use% Mounted on
/dev/sda2              303528624    6438700 281671544   3% /
udev                  1030332        116   1030216   1% /dev
/dev/sdb2              306128812    195568 290505224   1% /mnt

```

The running root file system device is the one mounted on /.

8. Dump the running root file system to the backup drive.

```

smw:~ # cd /mnt
smw:~ # dump 0f - \
/dev/disk/by-path/pci-0000:05:00.0-sas-phy7-0x4433221107000000-lun-0-part2 | restore rf -
DUMP: WARNING: no file '/etc/dumpdates'
DUMP: Date of this level 0 dump: Tue Mar 15 13:43:17 2011
DUMP: Dumping /dev/sda2 (/) to standard output
DUMP: Label: none
DUMP: Writing 10 Kilobyte records
DUMP: mapping (Pass I) [regular files]
DUMP: mapping (Pass II) [directories]
DUMP: estimated 7898711 blocks.
DUMP: Volume 1 started with block 1 at: Tue Mar 15 13:44:40 2011
DUMP: dumping (Pass III) [directories]
DUMP: dumping (Pass IV) [regular files]
restore: ./lost+found: File exists
DUMP: 79.34% done at 20890 kB/s, finished in 0:01
DUMP: Volume 1 completed at: Tue Mar 15 13:52:13 2011
DUMP: Volume 1 7908080 blocks (7722.73MB)
DUMP: Volume 1 took 0:07:33
DUMP: Volume 1 transfer rate: 17457 kB/s
DUMP: 7908080 blocks (7722.73MB)
DUMP: finished in 453 seconds, throughput 17457 kBytes/sec
DUMP: Date of this level 0 dump: Tue Mar 15 13:43:17 2011
DUMP: Date this dump completed: Tue Mar 15 13:52:13 2011
DUMP: Average transfer rate: 17457 kB/s
DUMP: DUMP IS DONE

```

9. Modify the backup drive's fstab and menu.lst files to reflect the backup drive's device, replacing the primary drive's device name.

NOTE: This step is necessary only if `by-id` names are used. If `by-path` names are being utilized for the root and swap devices, changes are not necessary; these devices reference physical slots,

and the backup drive will be moved to the same physical slot (slot 0) when replacing a failed primary boot drive.

- a. Edit `/mnt/etc/fstab`. Replace the `root` and `swap` partitions' `by-id` device names with those used for this backup device, replacing the original disk device name.

For example, change

```
/dev/disk/by-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RDS-part1 swap swap defaults
/dev/disk/by-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RDS-part2 / ext3 acl,user_xattr
```

to:

```
/dev/disk/by-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RD7-part1 swap swap defaults
/dev/disk/by-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RD7-part2 / ext3 acl,user_xattr
```

- b. Edit `/mnt/boot/grub/menu.lst`. Replace the `root=` and `resume=` device names with those used for this backup device, replacing the original disk device name.

The `root=` entry normally refers to partition `-part2`, and the `resume=` entry normally refers to partition `-part1`; these partition references must be maintained.

For example, replace the `menu.lst` configuration references of:

```
root=/dev/disk/by-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RDS-part2
```

with:

```
root=/dev/disk/by-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RD7-part2or
similarly with the by-id device names, if those are preferred.
```

Replace the `resume=` references similarly.

10. Install the `grub` boot loader. To make the backup drive bootable, reinstall the `grub` boot facility on that drive.



CAUTION: Although all of the disks connected to the SMW are available to the system, `grub` only detects the first 16 devices. Therefore, if you add a disk to the SMW after the SMW is connected to the boot RAID, it is advisable to reboot the SMW before continuing.

- a. Create a unique file on the backup drive to be used to identify that drive to `grub` boot facility.

```
smw:~ # cd /
smw:~ # touch /mnt/THIS_IS_6
```

- b. Invoke the `grub` boot utility. Within the `grub` boot utility:

1. Execute the `find` command to locate the drive designation that `grub` uses.
2. Select the drive to which the boot blocks will be installed with the `root` command.
3. Use the `setup` command to set up and install the `grub` boot blocks on that drive. The Linux `grub` utility and boot system always refer to drives as `hd`, regardless of the actual type of drives. For example:

```
smw:~ # grub --no-curses
GNU GRUB version 0.97 (640K lower / 3072K upper memory)
[ Minimal BASH-like line editing is supported. For the first word, TAB
lists possible command completions. Anywhere else TAB lists the
possible
completions of a device/filename. ]
```

```

grub> find /THIS_IS_6
(hd2,1)
grub> root (hd2,1)
root (hd2,1)
Filesystem type is ext2fs, partition type 0x83
grub> setup (hd2)
Checking if "/boot/grub/stage1" exists... yes
Checking if "/boot/grub/stage2" exists... yes
Checking if "/boot/grub/e2fs_stage1_5" exists... yes
Running "embed /boot/grub/e2fs_stage1_5 (hd2)"... 17 sectors are
embedded.
succeeded
Running "install /boot/grub/stage1 (hd2) (hd2)1+17 p (hd2,1)/boot/grub/
stage2 \
/boot/grub/menu.lst"... succeeded
Done.
grub> quit

```

IMPORTANT: For rack-mount SMWs, grub recreates device.map with the short names, not the persistent names. The /dev/sdx names must not be trusted. Always use find the next time you execute grub because it is possible that grub root may not be hd2 the next time you execute grub.

11. Unmount the backup root partition.

```
smw:~ # umount /mnt
```

The drive is now bootable once plugged in or cabled as the primary drive.

Desk-side SMW: Create an SMW Bootable Backup Drive

IMPORTANT: The disk device names shown in this procedure are only examples; substitute the actual disk device names for your system. For example, on an SMW with three SMW disks, the boot disk is /dev/sda and the bootable backup disk is /dev/sdc; on an SMW with two SMW disks, the boot disk is /dev/sda and the bootable backup disk is /dev/sdb.



CAUTION: Shut down the Cray system before you begin this procedure.

Also be aware that there may be a considerable load on the SMW while creating the SMW bootable backup drive.

1. Log on to the SMW as `crayadm` and `su` to `root`.

```
crayadm@smw:~> su - root
smw:~ #
```

2. If the backup drive disk partition table already exists and the partition table on the backup drive matches the partition table that is on the primary boot drive, skip this step; otherwise, create the backup drive disk partition table.

NOTE:

For optimal performance, the source and destination disks should be on different buses; drive slots 0 and 1 are on a different bus than drive slots 2 and 3.

In this example, the partition table consists of the following:

- Slice 1: 4 GB Linux swap partition
 - Slice 2: Balance of disk space used for the root file system
- a. Use the fdisk command to display the boot disk partition layout.

```
smw:~ # fdisk -lu /dev/sda
Disk /dev/sda: 320.0 GB, 320072933376 bytes
255 heads, 63 sectors/track, 38913 cylinders, total 625142448 sectors
Units = sectors of 1 * 512 = 512 bytes
```

Device	Boot	Start	End	Blocks	Id	System
/dev/sda1		63	8401994	4200966	82	Linux swap / Solaris
/dev/sda2	*	8401995	625137344	308367675	83	Linux

- b. Use the fdisk command to configure the bootable backup disk partition layout. Set the bootable backup disk partition layout to match the boot disk partition layout. First, clear all of the old partitions using the d command within fdisk; next create a Linux swap and a Linux partition; and then write your changes to the disk. For help, type m within fdisk (see the following sample output).

```
smw:~ # fdisk -u /dev/sdb

The number of cylinders for this disk is set to 38913.
There is nothing wrong with that, but this is larger than 1024,
and could in certain setups cause problems with:
1) software that runs at boot time (e.g., old versions of LILO)
2) booting and partitioning software from other OSs
   (e.g., DOS FDISK, OS/2 FDISK).

Command (m for help): p
Disk /dev/sdb: 320.0 GB, 320072933376 bytes
255 heads, 63 sectors/track, 38913 cylinders, total 625142448 sectors
Units = sectors of 1 * 512 = 512 bytes
```

Device	Boot	Start	End	Blocks	Id	System
/dev/sdb1		63	8401994	4200966	82	Linux swap
/dev/sdb2		8401995	625105214	308351610	83	Linux

```

Command (m for help): d
Partition number (1-5): 2
Command (m for help): d
Selected partition 1
Command (m for help): n
Command action
e extended
p primary partition (1-4)
p
Partition number (1-4): 1
First sector (63-625105215, default 63): (Press the Enter key)
Using default value 63
Last sector or +size or +sizeM or +sizeK (63-625105215, default 625105215):
8401994

Command (m for help): t
Selected partition 1
Hex code (type L to list codes): 82
Changed system type of partition 1 to 82 (Linux swap / Solaris)

Command (m for help): n
Command action
```

```

e extended
p primary partition (1-4)
p
Partition number (1-4): 2
First sector (8401995-625105215, default 8401995): (Press the Enter key)
Using default value 8401995
Last sector or +size or +sizeM or +sizeK (8401995-625105215, default
625105215): \
(Press the Enter key)
Using default value 625105215

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.

```

- c. Display the boot backup disk partition layout.

```

smw:~ # fdisk -lu /dev/sdb
Disk /dev/sdb: 320.0 GB, 320072933376 bytes
255 heads, 63 sectors/track, 38913 cylinders, total 625142448 sectors
Units = sectors of 1 * 512 = 512 bytes

```

Device	Boot	Start	End	Blocks	Id	System
/dev/sdc1		63	8401994	4200966	82	Linux swap / Solaris
/dev/sdc2	*	8401995	625137344	308367675	83	Linux

3. Initialize the swap device.

```

smw:~ # mkswap /dev/sdb1

```

4. (If the device names have already been standardized, skip this step.) Standardize the `/etc/fstab` and `grub` disk device names. The device names that the installation process writes into the `/boot/grub/menu.lst` file are UDEV-based names (for example, `/dev/disk/by-id/scsi-SATA_ST3320620AS_922J3-part2` or `/dev/disk/by-id/ata-ST3320620A_9QFA85PV-part2`) instead of the more commonly used device names (for example, `/dev/sda2` or `/dev/hda2`). In the following procedures, edit the `/boot/grub/menu.lst` file to change only the long UDEV-based name to the shorter, commonly used device name reflected in the output of the `df` command on your system. Be aware that errors in the `/boot/grub/menu.lst` will affect your ability to boot the SMW.

- a. SLES 11 sets up `/etc/fstab` and `/boot/grub/menu.lst` with UDEV-based names for the root device. For example:

```

smw:~ # head -2 /etc/fstab
/dev/disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part2 / ext3 acl,user_xattr 1
1
/dev/disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part1 swap swap defaults 0 0
smw:~ # more /boot/grub/menu.lst
###Don't change this comment - YaST2 identifier: Original name: linux###
title SUSE Linux Enterprise Server 11 - 2.6.27.19-5
root (hd0,1)
kernel /boot/vmlinuz-2.6.27.19-5-default \
root=/dev/disk/by-id/ata-ST3320620AS_5QF00F84-part2 \
resume=/dev/sda1 splash=silent crashkernel=256M-:128M@16M
showopts vga=0x31a initrd /boot/initrd-2.6.27.19-5-default

```

```
###Don't change this comment - YaST2 identifier: Original name: failsafe###
title Failsafe -- SUSE Linux Enterprise Server 11 - 2.6.27.19-5
root (hd0,1)
kernel /boot/vmlinuz-2.6.27.19-5-default \
root=/dev/disk/by-id/ata-ST3320620AS_5QF00F84-part2 showopts \
ide=nodma apm=off noresume edd=off powersaved=off nohz=off highres=off
processor.max_cstate=1 x11failsafe vga=0x31a
initrd /boot/initrd-2.6.27.19-5-default
```

- b. Execute the `df` command to get the name of the device to use in the `/etc/fstab` and `/boot/grub/menu.lst` files to replace the long UDEV-based device name.
For example:

```
smw:# df
Filesystem      1K-blocks    Used    Available    Use%    Mounted on
/dev/sda2        303528624  40652904   247457340    15%     /
udev             1030780     460       1030320      1%      /dev
```

- c. Create a backup copy of the `/etc/fstab` and `/boot/grub/menu.lst` files.

```
smw:# cp -p /etc/fstab /etc/fstab.save
smw:# cp -p /boot/grub/menu.lst /boot/grub/menu.lst.save
```

- d. Edit the `/etc/fstab` file appropriately, using the device name (`dev`) from the `df` command output. In this example, the "1" and "2" refer to the partition names on the device. Change the following lines, which changes the long name `disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part2` to `sda2` and changes `disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part1` to `sda1`. Ensure that your swap is on `sda1`:

```
/dev/disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part2 / ext3 acl,user_xattr 1
1
/dev/disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part1 swap swap defaults 0 0
to:
/dev/sda2 / ext3 acl,user_xattr 1 1
/dev/sda1 swap swap defaults 0 0
```

- e. Edit the `/boot/grub/menu.lst` file appropriately; use the device name (`dev`) from the `df` command output. Change the long name `disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part2` to `sda2`. Change the following lines:

```
title SUSE Linux Enterprise Server 11 - 2.6.27.19-5
kernel /boot/vmlinuz-2.6.27.19-5-default \
root=/dev/disk/by-id/ata-ST3320620AS_5QF00F84-part2 \
resume=/dev/sda1 splash=silent crashkernel=256M-:128M@16M showopts vga=0x31a
to
```

```
title SUSE Linux Enterprise Server 11 - 2.6.27.19-5
kernel /boot/vmlinuz-2.6.27.19-5-default \
root=/dev/sda2 resume=/dev/sda1 splash=silent \
crashkernel=256M-:128M@16M showopts vga=0x31a
```

and change the following lines:

```
title Failsafe -- SUSE Linux Enterprise Server 11 - 2.6.27.19-5
kernel /boot/vmlinuz-2.6.27.19-5-default \
root=/dev/disk/by-id/ata-ST3320620AS_5QF00F84part2 \
```

```
showopts ide=nodma apm=off noresume edd=off powersaved=off nohz=off \
highres=off processor.max_cstate=1 x11failsafe vga=0x31a
```

to

```
title Failsafe -- SUSE Linux Enterprise Server 11 - 2.6.27.19-5
kernel /boot/vmlinuz-2.6.27.19-5-default \
root=/dev/sda2 showopts ide=nodma apm=off noresume edd=off \
powersaved=off nohz=off highres=off processor.max_cstate=1 x11failsafe
vga=0x31a
```

- f. Verify that the edited files are correct and match the output of the `df` command.

```
smw:~ # head -2 /etc/fstab
/dev/sda2 / ext3 acl,user_xattr 1 1
/dev/sda1 swap swap defaults 0 0
smw:~ # more /boot/grub/menu.lst
###Don't change this comment - YaST2 identifier: Original name: linux###
title SUSE Linux Enterprise Server 11 - 2.6.27.19-5
root (hd0,1)
kernel /boot/vmlinuz-2.6.27.19-5-default root=/dev/sda2 \
resume=/dev/sda1 splash=silent crashkernel=256M-:128M@16M showopts
initrd /boot/initrd-2.6.27.19-5-default
###Don't change this comment - YaST2 identifier: Original name: failsafe###
title Failsafe -- SUSE Linux Enterprise Server 11 - 2.6.27.19-5
root (hd0,1)
kernel /boot/vmlinuz-2.6.27.19-5-default \
root=/dev/sda2 showopts ide=nodma apm=off noresume edd=off
powersaved=off nohz=off highres=off initrd /boot/initrd-2.6.27.19-5-default
```

5. Update the grub device table to recognize any new drives added since the initial operating system installation.

IMPORTANT: Although all of the disks connected to the SMW are available to the system, grub detects only the first 16 devices. Therefore, if you add a disk to the SMW after the SMW is connected to the boot RAID, it is advisable to reboot the SMW before continuing this procedure.

- a. Back up the current grub device mapping file.

```
smw:~ # mv /boot/grub/device.map /boot/grub/device.map-YYYYMMDD
```

- b. Invoke the grub utility to create a new device mapping file.

```
smw:~ # grub --device-map=/boot/grub/device.map
Probing devices to guess BIOS drives. This may take a long time.
GNU GRUB version 0.97 (640K lower / 3072K upper memory)
grub> quit
```

The file `/boot/grub/device.map` is now updated to reflect all drives, using the standardized drive naming. To verify the contents of the file:

```
smw:~ # cat /boot/grub/device.map
(fd0) /dev/fd0
(hd0) /dev/sda
(hd1) /dev/sdc
```

6. Create a new file system on the backup drive root partition by executing the `mkfs` command.

```
smw:~ # mkfs -t ext3 /dev/sdb2
mke2fs 1.41.1 (01-Sep-2008)
```

```

Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
19275776 inodes, 77091918 blocks
3854595 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=4294967296
2353 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
4096000, 7962624, 11239424, 20480000, 23887872, 71663616
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
This filesystem will be automatically checked every 33 mounts or
180 days, whichever comes first. Use tune2fs -c or -i to override.
smw:~ #

```

7. Mount the new backup root file system on /mnt.

```
smw:~ # mount /dev/sdb2 /mnt
```

8. Confirm the running root file system device.

```

smw:~ # df
Filesystem      1K-blocks    Used   Available   Use%    Mounted on
/dev/sda2       303528624  6438700   281671544    3%      /
udev            1030332     116     1030216     1%      /dev
/dev/sdb2       306128812  195568   290505224    1%      /mnt

```

The running root file system device is the one mounted on /.

9. Dump the running root file system to the backup drive.

```

smw:~ # cd /mnt
smw:~ # dump 0f - /dev/sda2 | restore rf -
DUMP: WARNING: no file `/etc/dumpdates'
DUMP: Date of this level 0 dump: Thu Nov 11 06:55:29 2010
DUMP: Dumping /dev/sda2 (/) to standard output
DUMP: Label: none
DUMP: Writing 10 Kilobyte records
DUMP: mapping (Pass I) [regular files]
DUMP: mapping (Pass II) [directories]
DUMP: estimated 4003398 blocks.
DUMP: Volume 1 started with block 1 at: Thu Nov 11 06:57:38 2010
DUMP: dumping (Pass III) [directories]
DUMP: dumping (Pass IV) [regular files]
restore: ./lost+found: File exists
DUMP: 81.99% done at 10941 kB/s, finished in 0:01
DUMP: Volume 1 completed at: Thu Nov 11 07:04:01 2010
DUMP: Volume 1 4008910 blocks (3914.95MB)
DUMP: Volume 1 took 0:06:23
DUMP: Volume 1 transfer rate: 10467 kB/s
DUMP: 4008910 blocks (3914.95MB)
DUMP: finished in 383 seconds, throughput 10467 kBytes/sec
DUMP: Date of this level 0 dump: Thu Nov 11 06:55:29 2010

```

```
DUMP: Date this dump completed: Thu Nov 11 07:04:01 2010
DUMP: Average transfer rate: 10467 kB/s
DUMP: DUMP IS DONE
```

10. To make the backup drive bootable, reinstall the grub boot facility on that drive.

- a. Create a unique file on the backup drive to be used to identify that drive to the grub boot facility.

```
smw:~ # cd /
smw:~ # touch /mnt/THIS_IS_SDx
```

- b. Invoke the grub boot utility. Within grub:

1. Execute the find command to locate the drive designation that grub uses.
2. Select the drive to which the boot blocks will be installed with the root command.
3. Use the setup command to set up and install the grub boot blocks on that drive.

```
smw:~ # grub --no-curses
GNU GRUB version 0.97 (640K lower / 3072K upper memory)
[ Minimal BASH-like line editing is supported. For the first word, TAB^[
lists possible command completions. Anywhere else TAB lists the possible
completions of a device/filename. ]
grub> find /THIS_IS_SDx
find /THIS_IS_SDx
(hd1,1)
grub> root (hd1,1)
root (hd1,1)
Filesystem type is ext2fs, partition type 0x83
grub> setup (hd1)
setup (hd1)
Checking if "/boot/grub/stage1" exists... yes
Checking if "/boot/grub/stage2" exists... yes
Checking if "/boot/grub/e2fs_stage1_5" exists... yes
Running "embed /boot/grub/e2fs_stage1_5 (hd1)"... 17 sectors are embedded.
succeeded
Running "install /boot/grub/stage1 (hd1) (hd1)1+17 p
(hd1,1)/boot/grub/stage2 /boot/grub/menu.lst"... succeeded
Done.
grub> quit
```

Note that the Linux grub utility and boot system always refer to drives as `hd`, regardless of the actual type of the drives.

11. Unmount the backup root partition.

```
smw:~ # umount /dev/sdb2
```

The drive is now bootable once plugged in or cabled as the primary drive.

Install an SMW Update Package



CAUTION: If you are updating a system set that is running, you must shut down your Cray system before installing the update package.

IMPORTANT: For a system configured for SMW high availability (HA) with the SMW failover feature, you must first prepare both SMWs for an upgrade.

Shut down the Cray system

1. Log on to the SMW as `crayadm` and confirm the Cray system is shut down.

```
crayadm@smw:~> ping boot
```

If the command responds with "alive", then it is up and needs to be shut down.

2. Shut down the system by typing the following command.

```
crayadm@smw:~> xtbootsys -s last -a auto.xtshutdown
```

For more information about using automation files, see the `xtbootsys(8)` man page.

Update the SMW software and configuration

1. Open a terminal window, and `su` to `root`.

```
crayadm@smw:~> su - root
smw:~ #
```

2. Mount the release media by using one of the following commands, depending on your media type.

If you are installing the update package from disk, place the `Cray SMW 7.2UPxx Software DVD` in the CD/DVD drive and mount it.

```
smw:~ # mount /dev/cdrom /media/cdrom
```

NOTE: If you want to install the update package from disk images instead of from the DVD, copy the files to a directory such as `/tmp/SMW_version` on the SMW and then substitute this path for `/media/cdrom` in subsequent instructions.

Or

If you are installing the update package using the ISO image, in the current directory execute the following command with the file name of the `smw-image` ISO image for the update that you are installing. For example:

```
smw:~ # mount -o loop,ro smw-image-7.2.0-1.0702.nnnnn.nnn-n.iso /media/cdrom
```

3. If you do not have `postfix` configured on your SMW, skip this step:

To prevent your `master.cf` and `main.cf` `postfix` configuration files from being recreated during software updates or fixes, ensure the following setting in your `/etc/sysconfig/mail` file on your SMW is set to "no":

```
MAIL_CREATE_CONFIG="no"
```

4. To see what SMW software will be updated in this new release, execute these commands prior to doing the update. This information is gathered, displayed, and contained in the log files during the `SMWinstall` process.

- a. Check security and recommended updates.

```
smw:~ # /media/cdrom/CRAYSMWinstall.sh -GS
```

- b. Check Cray software updates.

```
smw:~ # /media/cdrom/CRAYSMWinstall.sh -GV
```

5. Create a new copy of the SMWinstall.conf configuration file and modify your new copy of the SMWinstall.conf file for your site-specific requirements. You must be `root` to modify the SMWinstall.conf configuration file. The SMWinstall.conf configuration file is created during the installation process by copying the SMWinstall.conf template from the distribution media. By default, the SMW configuration file is placed in `/home/crayadm/SMWinstall.conf`.

```
smw:~ # cp /media/cdrom/SMWinstall.conf /home/crayadm
smw:~ # chmod 644 /home/crayadm/SMWinstall.conf
smw:~ # vi /home/crayadm/SMWinstall.conf
```

For a description of the contents of the SMWinstall.conf file, see the SMWinstall.conf(5) man page.

6. Update the software with SMWinstall. SMWinstall checks for any inconsistency between the system and the SMWinstall.conf file settings, prompts for the root MySQL database password, and stores its log files in `/var/adm/cray/logs`.

```
smw:~ # /media/cdrom/SMWinstall
...
```

```
Please enter your root DB password:
Please confirm your root DB password:
Password confirmed.
```

When SMWinstall finishes, it will suggest a reboot of the SMW.

7. If necessary, restore the locally modified versions of the following files.
 - a. If your site had locally modified versions of `/etc/syslog-ng/syslog-ng.conf` or `/etc/rsyslog.conf` before this SMW update, restore the local modifications to these files. During the upgrade procedure, the old files are saved in `/etc/syslog-ng/syslog-ng.conf-YYYYMMDDhhmm` and `/etc/rsyslog.conf.rpmsave`.
 - b. If your site had local modifications to `/opt/cray/hss/default/etc/auto.xtshutdown` before this SMW update, restore the local modifications to this file.
 - c. If your site had power management profiles, copy the files that you backed up from the `/opt/cray/hss/7.1.0/pm/profiles` directory into the new `/opt/cray/hss/default/pm/profiles` directory. See *Monitoring and Managing Power Consumption on the Cray XC30 System (S-0043)* for instructions on validating the restored profiles.
 - d. If your site had local modifications to `/opt/cray/hss/default/etc/sedc_srv.ini` before this SMW update, locate the destination of this symbolic link (see [Before You Start](#) on page 71), compare the content of the local file to the distributed version of the file, and update your local file appropriately.
8. If your site has an `/opt/cray/hss/default/etc/xtdiscover.ini` file, the SMW update process does not overwrite an existing `xtdiscover.ini` file; the new version is named `xtdiscover.ini.dist`. Compare the content of the new `xtdiscover.ini.dist` with your `xtdiscover.ini` file, and update your `xtdiscover.ini` file appropriately.

NOTE: If the `xtdiscover.ini` file does not exist, then the `xtdiscover.ini.dist` file is copied to the `xtdiscover.ini` file.

9. Unmount the SMW 7.2UPxx Software media.

```
smw:~ # umount /media/cdrom
```

10. If you are using the update disk, eject the SMW 7.2UPxx Software DVD.

```
smw:~ # eject
```

11. Reboot the SMW.

```
smw:~ # reboot
```

If you have a Cray XE or Cray XK system, continue to [For Cray XE or Cray XK Systems Only: Updating the L0 and L1 Firmware](#).

For Cray XC Series Systems Only: Update the BC and CC Firmware

Cray XC Series images for the cabinet controller (CC) and blade controller (BC) are always downloaded over the HSS network. Any updated firmware will be placed in `/opt/cray/hss-images/...` as part of the installation or update process. In order to boot the updated firmware, the `hss_make_default_initrd` script must be run, all CCs rebooted, and all BCs rebooted and power cycled.

IMPORTANT: If an installation step fails because of a hardware issue, such as a cabinet failing to power-up, when that issue is resolved you must go back to the last successful step in the installation procedure and continue from there. Do not skip steps or continue out of order.

For Cray XC Series systems only: Update the BC and CC firmware

Cray XC Series images for the cabinet controller (CC) and blade controller (BC) are always downloaded over the HSS network. Any updated firmware will be placed in `/opt/cray/hss-images/...` as part of the installation or update process. In order to boot the updated firmware, the `hss_make_default_initrd` script must be run, all CCs rebooted, and all BCs rebooted and power cycled.

IMPORTANT: If an installation step fails because of a hardware issue, such as a cabinet failing to powerup, when that issue is resolved you must go back to the last successful step in the installation procedure and continue from there. Do not skip steps or continue out of order.

1. Update the controller boot image.

The version used in the command argument for `hss_make_default_initrd` should match that of the version specified in the `lsb-cray-hss` line in the output from the `crms-release` file. This directory will not exist until the `hss_make_default_initrd` command completes.

```
smw:~ # cat /opt/cray/hss/default/etc/crms-release
HSS-CRMS Tue Oct 23 01:31:09 CDT 2012 on hssbld0 by bwdev
lsb-cray-hss-7.2.0-1.0702.32944.237
smw:~ # hss_make_default_initrd /opt/cray/hss-images/master/7.2.0-1.0702.32944.237
:: Verifying base RPM list to the manifest
:: Installing filesystem hierarchy
:: Installing busybox
:: Installing base RPMs
```

```

::: Installing ssh
::: Removing /etc/securetty installed by the pam RPM
::: Installing sshfs
::: Installing rsh
::: Modifying /etc/pam.d/rlogin to remove securetty checking
::: Modifying /etc/pam.d/rsh to remove rhosts and nologin checking (Bug #779466)
::: Installing rsync
::: Installing lsof
::: Installing atftp
::: Installing tcpdump
::: Installing ethtool
::: Installing syslog-ng
::: Installing logrotate
::: Installing ntp/ntpd
::: Installing strace
::: Installing screen
::: Installing minicom
::: Installing ppp
::: Installing mtd-utils
::: Installing /init
::: Installing file.rpm
::: Installing libgmodule
::: Installing Midnight Commander
::: Installing cray-viper
::: Installing spread
::: Installing coreboot-utils
::: Clearing init.d to be replaced by cray-hss32-filesystem
::: Creating initial etc files needed for root creation
::: Installing Cray kernel
::: Installing latest Cray kernel modules
::: Clearing select /boot items
::: Installing boot-parameters
::: Installing cray-hss32-scripts
::: Installing lsb-cray-hss-controllers
::: Installing cray-libconfig
::: Installing cray-bdm
::: Installing cray-play_xsvf
::: Removing unwanted files from the root

```

```

=====
The new initrd hierarchy is now in /opt/cray/hss-images/
master/ 7.2.0-1.0702.32944.237.

```

```

Running hssclone.
Image Clone Complete: /opt/cray/hss-images/default
Running hsspackage.
copying image
copying modules
running depmod
creating load file: /opt/cray/hss-images/default/HSS32.load
compressing initrd.img
Creating pxelinux.0 symlink
Running hssbootlink.
linking /opt/cray/hss-images/default/HSS32/bzImage-3.0.13-0.2 7.2_1.0700.6694-
cray_hss32 \
/opt/tftpboot/bzImage
linking /opt/cray/hss-images/default/HSS32/parameters \
/opt/tftpboot/pxelinux.cfg/default
linking /opt/cray/hss-images/default/HSS32/initrd.img /opt/tftpboot/initrd.img

```

2. Power down the system.

```
smw:~ # xtcli power down s0
```

3. Reboot the cabinet controllers, then ensure that all cabinet controllers are up.

```
smw:~ # xtccreboot -c all
xtccreboot: reboot sent to specified CCs
smw:~ # xtalive -l cc
```

4. Power up the system.

```
smw:~ # xtcli power up s0
```

Note that at this point the `xtcli` status output shows that all nodes are "off", because they have not yet been bounced.

5. Run the `xtdiscover` command to ensure that any changes made to the HSS database schema for new features are captured.

```
smw:~ # xtdiscover
```

6. Exit from the `root` login.

```
smw:~ # exit
```

7. Run the `rtr --discover` command to determine the exact configuration of the HSN.

```
crayadm@smw:~> rtr --discover
```

If the system was not bounced previously you may see the following message:

```
System was not bounced in diagnostic mode, should I re-bounce? Continue (y/n)?
```

If so, respond with **y**.

8. Update the firmware. Execute the `xtzap` command to update the components.



CAUTION: The `xtzap` command is normally intended for use by Cray Service personnel only. Improper use of this restricted command can cause serious damage to your computer system.

```
crayadm@smw:~> xtzap -r -v s0
```

IMPORTANT: The Cray XC30 system also requires an update to the NVIDIA® BIOS (nvBIOS) for the NVIDIA K20X graphics processing units (GPUs). This update is done after CLE has been booted. For more information, see *Installing and Configuring Cray Linux Environment (CLE) Software (S-2444)*.

9. Use the output of the `xtzap` command to determine if any components need to be flashed.

While the `xtzap -a` command can be used to update all components with a single command, it may be faster to use the `xtzap -blade` command when only blade types need to be updated, or the `xtzap -t` command when only a single type needs to be updated. On larger systems, this can be a significant time savings.

This is the list of all cabinet level components:

`cc_mc` (CC Microcontroller)

cc_bios (CC Tolapai BIOS)
 cc_fpga (CC FPGA)
 chia_fpga (CHIA FPGA)

This is a list of all blade level components:

cbb_mc (CBB BC Microcontroller)
 ibb_mc (IBB BC Microcontroller)
 anc_mc (ANC BC Microcontroller)
 bc_bios (BC Tolapai BIOS)
 lod_fpga (LOD FPGA)
 node_bios (Node BIOS)
 loc_fpga (LOC FPGA)
 qloc_fpga (QLOC FPGA)

If the output of the `xtzap` command shows that only a specific type needs to be updated, then use the `-t` option with that type (this example uses the `node_bios` type).

```
crayadm@smw:~> xtzap -t node_bios s0
```

If the output of the `xtzap` command shows that only blade component types need to be updated, then use the `-b` option:

```
crayadm@smw:~> xtzap -b s0
```

If the output of the `xtzap` command shows that both blade- and cabinet-level component types need to be updated, or if you are unsure what needs to be updated, then use the `-a` option:

```
crayadm@smw:~> xtzap -a s0
```

10. Execute the `xtzap -r -v s0` command again; all firmware revisions should report correctly, except `node_bios`; `node_bios` will display as "NOT_FOUND" until after you run the `xtbounce --linktune` command.

```
crayadm@smw:~> xtzap -r -v s0
```

11. Execute the `xtbounce --linktune` command, which forces `xtbounce` to do full tuning on the system.

```
crayadm@smw:~> xtbounce --linktune=all s0
```

Continue with [Confirm the SMW Is Communicating with the System Hardware](#).

For Cray XE or Cray XK Systems Only: Verify Hardware Programmable Interface Controllers (PICs)

Verify Hardware Programmable Interface Controllers (PICs)

Execute the `xtcheckpic` command to verify the PICs are current. If any PICs need to be updated, contact your local Cray Service personnel. The PIC codes will be updated according to the site staff's priority; the PICs

need not be updated while performing this SMW update. Be aware, however, that updating PICs at a later date will require dedicated system time.

```
crayadm@smw:~> xtcheckpic
Checking PIC revisions, please wait. This could take a few minutes... Done.
```

```
*****
*      Current PIC Revisions      *
*****
* XT3-XT5 Verty          == 0x0d *
* Seastar Mezzanine PIC  == 0x15 *
* XT3 940 CPU PIC        == 0x1e *
* XT3 SIO PCIX Riser PIC == 0x1b *
* XT4 AM2 CPU PIC        == 0x0d *
* XT3 SIO PCIE Riser PIC == 0x07 *
* XR1 Riser PIC          == 0x03 *
* XT5 CPU PIC            == 0x13 *
* XT5 VRM PIC on XT5     == 0x06 *
* Crayfish Mezzanine PIC == 0x02 *
* Gemini Mezzanine PIC    == 0x17 *
* XIO CPU PIC            == 0x09 *
* XT5 VRM PIC on XIO     == 0x07 *
* XIO PCIE PIC           == 0x04 *
* G34 VERTY              == 0x09 *
* G34 CPU PIC            == 0x0e *
* G34 OLD! VRM PIC on G34 == 0x01 *
* G34 VRM PIC ON G34     == 0x01 *
* XMT2 CPU PIC           == 0x03 *
* XMT2 VRM on XMT2       == 0x01 *
* BAX CPU PIC            == 0x06 *
* G34 VRM PIC on BAX     == 0x01 *
* BAX PCI-E PIC          == 0x08 *
* NP VERTY               == 0x02 *
*****

!!INCORRECT BAX CPU PIC      : c0-0c2s2 0:0x60      version=0x05
!!INCORRECT BAX CPU PIC      : c0-0c2s2 1:0x60      version=0x05
!!INCORRECT BAX CPU PIC      : c0-0c2s2 2:0x60      version=0x05
!!INCORRECT BAX CPU PIC      : c0-0c2s2 3:0x60      version=0x05
```

For Cray Systems with Cray XK Blades Only: Update the Cray XK Accelerator Firmware

IMPORTANT: For Cray systems with Cray XK blades: After installing the SMW software update package, you must contact your Cray Service personnel to update the Cray XK accelerator firmware so that it is synchronized with the CLE release of the Cray XK driver. These steps must be done before installing CLE or booting the Cray system and can be done even if your site does not install the latest CLE release package.

Update SMW Software on the Boot Root and Shared Root

The SMWinstallCLE script updates the SMW software on the boot root and shared root for systems already running the Cray Linux Environment (CLE) software. For more information about the SMWinstallCLE script, see the SMWinstallCLE(8) man page.

Update SMW software on the boot root and shared root

NOTE: The RPMs that are installed on the boot root and shared root by `SMWinstallCLE` will also be installed the next time that `CLEinstall` is executed if the RPMs were not installed by `SMWinstallCLE`. Use `SMWinstallCLE` if you want to boot CLE after the SMW update. However, if you will be updating CLE software without booting CLE after the SMW update, then this procedure may be skipped.

1. As `root`, mount the release media by using one of the following commands, depending on your media type.

If you are installing the update package from disk, place the `SMW 7.2UPxx Software DVD` in the CD/DVD drive and mount it.

```
smw:~ # mount /dev/cdrom /media/cdrom
```

NOTE: If you want to install the update package from disk images instead of from the DVD, copy the files to a directory such as `/tmp/SMW_version` on the SMW and then substitute this path for `/media/cdrom` in subsequent instructions.

Or

If you are installing the update package using the ISO image, in the current directory execute the following command with the file name of the `smw-image` ISO image for the update that you are installing. For example:

```
smw:~ # mount -o loop,ro smw-image-7.2.0-1.0702.nnnnn.nnn-n.iso /media/cdrom
```

2. In the following steps it is assumed that the label `label_name` is described in the `/etc/sysset.conf` system set configuration file. The following example updates the `label_name` system set from the `/etc/sysset.conf` system set configuration file. See the `sysset.conf(5)` man page for additional information about the `/etc/sysset.conf` file.

```
smw:~ # /media/cdrom/utils/SMWinstallCLE --label=label_name
```

NOTE: The `SMWinstallCLE` command checks whether the boot node is booted. If it is booted, `SMWinstallCLE` prompts you to confirm that the system set you are changing is not the one booted.

```
HH:MM:SS WARNING: Your bootnode is booted. Please confirm that the system
set you
intend to update is not booted.
Do you wish to proceed? [n] y
```

For more detailed information, see the `SMWinstallCLE(8)` man page.

3. Unmount the `SMW 7.2UPxx Software` media.

```
smw:~ # umount /media/cdrom
```

4. If you are using the update disk, eject the `SMW 7.2UPxx Software DVD`.

```
smw:~ # eject
```

Your SMW software is now updated. If you have yet to configure your firewall, see [Set up the SUSE firewall and IP tables](#) on page 32. Then continue to install your CLE software using *Installing and Configuring Cray Linux Environment (CLE) Software (S-2444)*, which is provided with your CLE release package.

NOTE: To reconfigure a `LOGDISK`, `PMDISK`, or `DBDISK` when you need to replace a failed drive with a new drive on a rack-mount SMW, see [Rack-mount SMW: Replace a failed disk drive](#) on page 109.

Configure the Simple Event Correlator (SEC)

The System Management Workstation (SMW) 7.2.UP00 release includes the Open Source simple event correlator (SEC) package, `sec-2.7.0`, and an SEC support package, `cray-sec-version`. The SEC support package contains control scripts to manage the starting and stopping of SEC around a Cray mainframe boot session, in addition to other utilities.

If you plan to use the Cray SEC, see *Configuring SEC Software for a Cray XC, Cray XE, or Cray XK System* (S-2542) for configuration procedures.

Install Additional Software as RPMs

A variety of software packages are distributed as standard Linux RPM Package Manager (RPM) packages. RPM packages are self-contained installation files that create all required directories and install all component files in the correct locations when executed with the `rpm` command.

To install RPMs on the SMW, as `root`:

```
smw:~ # rpm -ivh /directorypath/filename.rpm
```

For more information, see the `rpm(8)` man page.

Enable Remote Access to the SMW

Virtual Network Computing (VNC) software enables you to view and interact with the SMW from another computer. The Cray system provides a VNC server, `Xvnc` and a VNC account `cray-vnc`. You must download a VNC client to connect to it. See RealVNC (<http://www.realvnc.com/>) or TightVNC (<http://www.tightvnc.com/>) for more information.

The VNC software requires a TCP/IP connection between the server and the viewer. Be aware that VNC is considered to be an insecure protocol, therefore some firewall security policies may not allow this connection. To create a secure VNC connection, see the topics on creating connecting to the VNC server through an SSH tunnel.

Start the VNC server

1. Log on to the SMW as `root` user.
2. Use the `chkconfig` command to check the current status of the server:

```
smw:~ # chkconfig vnc
vnc    off
```

3. If the `chkconfig` command you executed in step 2 reports that `Xvnc` was started by INET services (`xinetd`) execute the following commands to disable `xinetd` startup of `Xvnc`. Startup of `Xvnc` via `xinetd` is the SLES 11 default, but it usually is disabled by `chkconfig`:

```
smw:~ # chkconfig vnc off
smw:~ # /etc/init.d/xinetd reload
Reload INET services (xinetd).                done
```

If no other `xinetd` services have been enabled, the `reload` command will return `failed` instead of `done`. If the `reload` command returns `failed`, this is normal and you can ignore the notification.

4. Use the `chkconfig` command to start `Xvnc` at boot time:

```
smw:~ # chkconfig vnc on
```

5. Start the `Xvnc` server immediately:

```
smw:~ # /etc/init.d/vnc start
```

If the password for `cray-vnc` has not already been established, the system prompts you for one. You must enter a password to access the server.

```

Password: *****
Verify:
Would you like to enter a view-only password (y/n)? n
xauth:  creating new authority file /home/cray-vnc/.Xauthority

New 'X' desktop is smw:1

Creating default startup script /home/cray-vnc/.vnc/xstartup
Starting applications specified in /home/cray-vnc/.vnc/xstartup
Log file is /home/cray-vnc/.vnc/smw:1.log

```

To access the Xvnc server, use a VNC client, such as `vncviewer`, `tight_VNC`, `vnc4`, or a web browser. Direct it to the SMW that is running Xvnc. Many clients allow you to specify whether you want to connect in view-only or in an active mode. If you choose active participation, every mouse movement and keystroke made in your client is sent to the server. If more than one client is active at the same time, your typing and mouse movements are intermixed.

Commands entered through the VNC client affect the system as if they were entered from the SMW. However, the main SMW window and the VNC clients cannot detect each other. It is a good idea for the administrator who is sitting at the SMW to access the system through a VNC client.

The startup script starts the Xvnc server for display :1.

6. Verify that Xvnc started:

```

smw:~ # ps -e | grep vnc
1839 pts/0    00:00:00 Xvnc

```

For workstation or laptop running Linux: Connecting to the VNC server through an SSH tunnel, using the `vncviewer -via` option

IMPORTANT: This procedure is for use with the TightVNC client program.

Verify that you have the `vncviewer -via` option available. If you do not, use [For workstation or laptop running Linux: Connecting to the VNC server through an SSH tunnel](#) on page 99.

If you are connecting from a workstation or laptop running Linux, enter the `vncviewer` command shown below.

The first password you enter is for `crayadm` on the SMW. The second password you enter is for the VNC server on the SMW, which was set first time the VNC server was started using `/etc/init.d/vnc start`.

```

/home/mary> vncviewer -via crayadm@smw localhost:1
Password: *****
VNC server supports protocol version 3.130 (viewer 3.3)
Password: *****
VNC authentication succeeded
Desktop name "cray-vnc's X desktop (smw:1)"
Connected to VNC server, using protocol version 3.3

```

For workstation or laptop running Linux: Connecting to the VNC server through an SSH tunnel

NOTE: This procedure assumes that the VNC server on the SMW is running with the default port of 5901.

1. This `ssh` command starts an `ssh` session between the local Linux computer and the SMW, and it also creates an SSH tunnel so that port 5902 on the `localhost` is forwarded through the encrypted SSH tunnel to port 5901 on the SMW. You will be prompted for the `crayadm` password on the SMW.

```
local_linux_prompt> ssh -L 5902:localhost:5901 smw -l crayadm
Password:
crayadm@smw>
```

2. Now `vncviewer` can be started using the local side of the SSH tunnel, which is port 5902. You will be prompted for the password of the VNC server on the SMW. This password was set when the VNC server was started for the first time using `/etc/init.d/vnc start` on the SMW.

```
local_linux_prompt> vncviewer localhost:2
Connected to RFB server, using protocol version 3.7
Performing standard VNC authentication
Password:
```

The VNC window from the SMW appears. All traffic between the `vncviewer` on the local Linux computer and the VNC server on the SMW is now encrypted through the SSH tunnel.

For workstation or laptop running Mac OS X: Connecting to the VNC server through an SSH tunnel

NOTE: This procedure assumes that the VNC server on the SMW is running with the default port of 5901.

1. This `ssh` command starts an `ssh` session between the local Mac OS X computer and the SMW, and it also creates an SSH tunnel so that port 5902 on the `localhost` is forwarded through the encrypted SSH tunnel to port 5901 on the SMW. You will be prompted for the `crayadm` password on the SMW.

```
local_mac_prompt> ssh -L 5902:localhost:5901 smw -l crayadm
Password:
crayadm@smw>
```

2. Now `vncviewer` can be started using the local side of the SSH tunnel, which is port 5902. You will be prompted for the password of the VNC server on the SMW. This password was set when the VNC server was started for the first time using `/etc/init.d/vnc start` on the SMW.

The SSH tunnel is now ready for use. To bring up the `vncviewer` on a Mac OS X computer, type the following on the command line:

```
local_mac_prompt% open vnc://localhost:5902
```

All traffic between the `vncviewer` on the local Mac computer and the VNC server on the SMW is encrypted through the SSH tunnel.

For workstation or laptop running Windows: Connecting to the VNC server through an SSH tunnel

NOTE: If you are connecting from a computer running Windows, then both a VNC client program, such as TightVNC and an SSH program, such as PuTTY, SecureCRT, or OpenSSH are recommended.

1. The same method described in [For workstation or laptop running Linux: Connecting to the VNC server through an SSH tunnel](#) can be used for computers running the Windows operating system.

Although TightVNC encrypts VNC passwords sent over the network, the rest of the traffic is sent unencrypted. To avoid a security risk, install and configure an SSH program that creates an SSH tunnel between TightVNC on the local computer (localhost port 5902) and the remote VNC server (localhost port 5901).

NOTE: Because the procedure for creating the SSH tunnel varies among different SSH programs for Windows computers, no instructions are provided here.

2. After installing TightVNC, start the VNC viewer program by double-clicking on the TightVNC icon. Enter the hostname and VNC screen number, `localhost: number` (such as, `localhost:2` or `localhost:5902`), and then select the Connect button.

Set up the SUSE firewall and IP tables

The SMW software includes a firewall. The following steps are for enabling and configuring the firewall on your system.

NOTE: You do not have to shut down the system to perform this task.

1. Before modifying the SUSE Firewall settings, make a copy of the configuration file:

```
smw:~ # cp -p /etc/sysconfig/SuSEfirewall12 /etc/sysconfig/SuSEfirewall12.orig
```

2. By using the SuSEfirewall12 program and the following procedure, change the IP tables rules to close off all unnecessary ports on the SMW.

```
smw:~ # iptables -L
smw:~ # vi /etc/sysconfig/SuSEfirewall12
```

Change the settings of these variables to the values shown:

```
FW_DEV_EXT="any eth0"

FW_DEV_INT="eth1 eth2 eth3 eth4 lo"

FW_SERVICES_EXT_UDP="161"

FW_TRUSTED_NETS="your_bootnode_ipaddress,tcp,7004 \
    your_syslognode_ip,udp,514 your_sdbnode_ip,tcp,6811:6815"
```

For example:

```
smw:~ # diff /etc/sysconfig/SuSEfirewall12.orig /etc/sysconfig/SuSEfirewall12
99c99
< FW_DEV_EXT="eth-id-00:30:48:5c:b0:ee eth0"
---
> FW_DEV_EXT="any eth0"
114c114
< FW_DEV_INT="eth-id-00:0e:0c:b4:df:64 eth-id-00:0e:0c:b4:df:65
    eth-id-00:0e:0c:b4:df:66 eth-id-00:0e:0c:b4:df:67 eth1 eth2 eth3 eth4"
---
> FW_DEV_INT="eth1 eth2 eth3 eth4 lo"
263c263
< FW_SERVICES_EXT_UDP=""
---
> FW_SERVICES_EXT_UDP="161"
394c394
< FW_TRUSTED_NETS=""
---
> FW_TRUSTED_NETS="10.3.1.254,tcp,7004 10.5.1.2,udp,514 10.5.1.2,tcp,6811:6815"
```

3. Invoke the modified configuration.

```
smw:~ # /etc/init.d/SuSEfirewall12_init start
smw:~ # /etc/init.d/SuSEfirewall12_setup start
```

4. Execute the following commands to start the firewall at boot time.

```
smw:~ # chkconfig SuSEfirewall12_init on
smw:~ # chkconfig SuSEfirewall12_setup on
```

5. To verify your changes to the iptables:

```
smw:~ # iptables -nvL
```

SSH access is one of the protocols permitted through the firewall from the external network to the SMW. For information about how to use Virtual Network Computing (VNC) through an SSH tunnel, see [Enable Remote Access to the SMW](#) on page 97.

Enable an Integrated Dell Remote Access Controller (iDRAC6) on a rack-mount SMW

Prerequisites

To enable an iDRAC6 on a rack-mount SMW, you need:

- Physical access to the SMW console
- The iDRAC6 IP address, subnet mask, and default gateway
- The SMW root account password

TIP: The following command retrieves the iDRAC IP information:

```
smw:~> /usr/bin/ipmitool -I open lan print 1
```

1. If the SMW is running, `su` to `root` and shut it down.

```
crayadm@smw:~> su - root
smw:~ # shutdown -h now;exit
```

2. Connect Ethernet cable to the iDRAC6 port. The cable is located on back of a rack-mount SMW in the lower left corner.
3. Power up the SMW.
4. After the BIOS, Dell PowerEdge Expandable RAID Controller (PERC) card, and disk map have displayed, the IPv4/IPv6 information displays. When the IPv4/IPv6 information displays, press **Ctrl-E**.
5. Using the arrow keys, select **LAN Parameters**, then press **Enter**.
6. Select **NIC Selection** and set it to **Dedicated**. Then press **Esc**.
7. Using the arrow keys, scroll down and select the **IPv4 settings** section.
 - a. Ensure that IPv4 is enabled.
 - b. Confirm that the IPv4 address source is set to static:

```
IPv4 Address Source: Static
```

- c. Enter your iDRAC6 IP addresses for the following:
 - Address:
 - Subnet Mask:
 - Default Gateway:

-
- d. Ensure that IPv6 is disabled.
 - e. Press **Esc** and return to the **LAN Parameters** window.
8. Using the arrow keys, select **LAN User Configuration**, then press **Enter**.
- NOTE:** This configuration is for both SSH and web browser access to the iDRAC.
- a. Enter the `root` account name and iDRAC password.
- ```
Account User name: root
Enter Password: *****
Confirm Password: *****
```
- a. Press **Esc**.
  - a.
9. Press **Esc** again.
10. Select **Save Changes and Exit**, then press **Enter**.

The SMW will complete booting up; no user interaction is required.

---

## Change the default iDRAC Password

---

1. Log into the web interface as `root`.
2. Select iDRAC settings on the left-hand bar.
3. Select network/Security on the main top bar.
4. Select Users on the secondary top bar.
5. Select the user whose password you are changing. For example, `userid 2` and `username root`.
6. Select Configure User, then Next.
7. Type the new password into the New Password and Confirm New Password text boxes.
8. Select Apply to complete the password change.

## Using the iDRAC6

---

1. Bring up a web browser.
2. Go to: `https://cray-drac`. A login screen appears.
3. Enter the account user name and password that you set up in [Enable an Integrated Dell Remote Access Controller \(iDRAC6\) on a rack-mount SMW](#) on page 103, then select **Submit**.  
The **System Summary** window appears.
4. To access the SMW console, select the **Console Media** tab.  
The **Virtual Console and Virtual Media** window appears.
5. Select **Launch Virtual Console**.

**TIP:**

By default, your console window has two cursors: one for the console and one for your own window environment. To switch to single-cursor mode, select **Tools**, then **Single Cursor**. This single cursor will not move outside the console window. To exit single-cursor mode, press the **F9** key.

To log out of the virtual console, kill the window or select **File**, then **Exit**. You will still be logged into the iDRAC6 in your web browser.

For detailed information, download the iDRAC6 documentation at:

<http://goo.gl/4Jm4T>

---

## Rack-mount SMW: Install the Dell Systems Management Tools and Documentation DVD

---

For advanced control over the Integrated Dell Remote Access Controller (iDRAC) and to provide features such as Automatic Recovery (automatic system boot after a power event), install the Server Administrator software from the Dell Systems Management Tools and Documentation DVD. If you have a Rack-mount SMW, this DVD was provided with your system.

1. Obtain the Dell System Management Tools and Documentation DVD.
2. Login as `root` on the SMW.
3. Mount the DVD.

```
smw:~ # mount /dev/cdrom /media/cdrom
```

4. Go to the location of the installation scripts.

```
smw:~ # cd /media/cdrom/SYSMGMT/srvadmin/linux/supportscripts
```

5. Execute the script to install the software.

```
smw:~ # sh srvadmin-install.sh --express
```

6. Start the Server Administrator services.

```
smw:~ # sh srvadmin-services.sh start
```

7. Double-click the icon named Launch Server Administrator on the SMW screen.
8. Enter the SMW user name `root`.
9. Enter the SMW `root` account password.

Now you are ready to manage the system for Properties, Shutdown, Logs, Alert Management, and Session Management.

The following websites may be useful:

- Dell OpenManage Linux Repository - See <http://linux.dell.com/wiki/index.php/Repository/OMSA>
- Dell OpenManage Server Administrator Version 7.0 Manuals (including complete installation guide and user guide - See <http://support.dell.com/support/edocs/software/srvadmin/7.0/en/index.htm>

---

## Rack-mount SMW: Configure the SMW to restore AC power after a power failure

---

1. Double-click the icon named Launch Server Administrator on the SMW screen to bring up the GUI.
2. Enter the SMW `root` user name and password.
3. In the left pane, select System.
4. In the left pane, select Main System Chassis.
5. In the left pane, select BIOS.
6. In the middle pane, select Setup Tab.
7. Select AC Power Recovery and select the value field.
8. Check Last.
9. Select Apply Change.
10. Select Go Back To BIOS Page.
11. Select OS Watchdog Timer, which should be set to Disable.

If the OS Watchdog Timer is not set to Disable, select the value field, check Disable, and then select Apply Change.

12. Exit the GUI.

Rack-mount SMW: Replace a Failed `LOGDISK` or `DBDISK` Disk Drive

## Rack-mount SMW: Replace a failed disk drive

The following steps are for replacing a failed LOGDISK, DBDISK, or PMDISK. You must be running the SUSE Linux Enterprise Server version 11 Service Pack 3 (SLES 11 SP3) SMW base operating system and a release of SMW 7.2 or later on your SMW in order to perform this task.

**TIP:** These instructions specify replacing the LOGDISK disk. If you are replacing the DBDISK or PMDISK, use the appropriate /dev/disk/by-path/pci\* device name.

1. Replace the failed drive with the new drive.
2. Reboot the SMW.

```
smw:~ # reboot
```

3. Reconfigure LOGDISK.

```
smw:~ # /sbin/fdisk /dev/disk/by-path/pci-0000:05:00.0\
-sas-phy4-0x4433221104000000-lun-0
Command (m for help): p
```

- a. Delete all the current partitions, if there are any.

```
Command (m for help): d
Partition number 4
Command (m for help): d
Partition number 3
Command (m for help): d
Partition number 2
Command (m for help): d
Partition number 1
Command (m for help): p
```

- b. Create the new, single partition.

```
Command (m for help): n
Command action
 e extended
 p primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-121601, default 1): # Hit return, take the default
Using default value 1
Last cylinder, +cylinders or +size{K,M,G} (1-121601, default 121601):
Hitreturn, take the default
Using default value 121601

Command (m for help): p
```

```

Disk /dev/disk/by-path/pci-0000:05:00.0-sas-phy4-0x4433221104000000-
lun-0:1000.2 GB, \
1000204886016 bytes
255 heads, 63 sectors/track, 121601 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Disk identifier: 0x00000083

Device Boot
Start End Blocks Id System
/dev/disk/by-path/pci-0000:05:00.0-sas-phy4-0x4433221104000000-lun-0-part1 1
121601 976760001 83 Linux

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.

```

4. Recreate the file system.

```

smw:~ # mkfs -t ext3 -b 4096
/dev/disk/by-path/pci-0000:05:00.0-sas-phy4-0x4433221104000000-lun-0-part1

```

5. Mount the newly created file system.

```

smw:~ # mount /dev/disk/by-path/pci-0000:05:00.0-sas-phy4-0x4433221104000000\
-lun-0-part1/ var/opt/cray/disk/1

```

6. The symbolic links should already be there to link this to /var/opt/cray/log.

```

smw:~ # ls -al /var/opt/cray

```

7. Create the following new directories:

```

smw:~> mkdir /var/opt/cray/disk/1/log
smw:~> mkdir /var/opt/cray/disk/1/debug
smw:~> mkdir /var/opt/cray/disk/1/dump

```

8. Restart the `rsms` daemon.

```

smw:~ # /etc/init.d/rsms restart
smw:~ # /etc/init.d/dbMonitor restart

```