



System Management Workstation (SMW) Software Installation Guide S-2480-7204a

Contents

About the System Management Workstation (SMW) Software Installation Guide.....	5
Related Publications.....	6
Distribution Media.....	6
Install the SMW Release Package.....	7
Prepare to Install the SMW Release Package.....	9
Network Connections.....	9
Configuration Values.....	9
Passwords.....	11
Rack-mount SMW: Linux /dev Names Mapped to Physical Drive Slots.....	11
Install the Base Operating System.....	13
R630 SMW: Configure the RAID Virtual Disks as Non-RAID.....	14
R630 SMW: Change the BIOS and iDRAC Settings.....	16
R630 SMW: Install the SLES Base Operating System.....	26
R815 SMW: Enable an Integrated Dell Remote Access Controller (iDRAC).....	31
R815 SMW: Install the SLES Base Operating System.....	36
Desk-side SMW: Install the SLES Base Operating System.....	41
Install and Configure SMW Software	45
Discover the Cray System Hardware and Power Up the Full System.....	48
Bootstrap Hardware Discovery.....	48
Discover the Cray System Hardware.....	52
Confirm the SMW is Communicating with System Hardware.....	58
Change the Default System Management Workstation (SMW) Passwords.....	60
Change the Default iDRAC Password.....	60
Set Up the SUSE Firewall and IP Tables.....	61
Configure the Boot RAID.....	62
Configure and Zone the Boot RAID for a DDN Storage System.....	63
Configure the Boot RAID for a NetApp, Inc. Engenio Storage System.....	71
Zone the QLogic™ FC Switch.....	75
Reboot the SMW and Verify LUNs are Recognized.....	78
Partition the LUNs.....	78
Set Up Boot RAID Space for Direct-attached Lustre.....	78
Configure the Simple Event Correlator (SEC).....	79
Back Up the Newly Installed and Configured SMW Software.....	80
R630 SMW: Create an SMW Bootable Backup Drive.....	80
(Optional) R630 SMW: Set Up the Bootable Backup Drive as an Alternate Boot Device.....	88

R815 SMW: Create an SMW Bootable Backup Drive.....	90
(Optional) R815 SMW: Set Up the Bootable Backup Drive as an Alternate Boot Device.....	98
Desk-side SMW: Create an SMW Bootable Backup Drive.....	100
(Optional) Desk-side SMW: Set Up the Bootable Backup Drive as an Alternate Boot Device.....	106
Update or Upgrade the Cray SMW Software.....	109
Prepare to Upgrade or Update SMW Software.....	109
Update the properties.local File.....	110
Back Up the Current Software.....	110
Install an SMW Update Package.....	132
Shut Down the Cray System.....	132
Update the SMW Software and Configuration.....	133
For Cray XC Series Systems Only: Update the BC and CC Firmware.....	135
Update the L0 and L1 Firmware.....	138
For Cray XE or Cray XK Systems Only: Verify Hardware Programmable Interface Controllers (PIC).....	139
For Cray Systems with Cray XK Blades Only: Update the Cray XK Accelerator Firmware.....	140
Update SMW Software on the Boot Root and Shared Root.....	140
Set Up the SUSE Firewall and IP Tables.....	141
Configure the Simple Event Correlator (SEC).....	142
Perform Additional SMW Installation or Configuration Procedures.....	143
Install Additional Software as RPMs.....	143
Use the iDRAC.....	143
Change the Time Zone for the SMW and the Blade and Cabinet Controllers.....	144
Rack-mount SMW: Install the Dell Systems Management Tools and Documentation DVD.....	147
Enable Remote Access to the SMW using VNC.....	147
Start the VNC server.....	148
For Workstation or Laptop Running Linux: Connecting to the VNC Server through an SSH Tunnel, using the <code>vncviewer -via</code> Option.....	149
For Workstation or Laptop Running Linux: Connecting to the VNC Server through an SSH Tunnel.....	150
For Workstation or Laptop Running Mac OS X: Connecting to the VNC Server through an SSH Tunnel.....	150
For Workstation or Laptop Running Windows: Connecting to the VNC Server through an SSH Tunnel.....	151
Troubleshoot SMW Problems.....	152
R630 SMW: System Repair.....	152
R815 SMW: Configure the SMW to Restore AC Power After a Power Failure	152
Replace a Failed Disk Drive.....	153
R630 SMW: Replace a Failed LOGDISK.....	153

R815 SMW: Replace a Failed LOGDISK.....	155
Recover from a Corrupt or Missing HSS MySQL Database.....	157
Restore a Corrupt or Missing HSS MySQL Database from a Backup.....	157
Regenerate a Corrupt or Missing HSS MySQL Database from Scratch.....	158
Troubleshoot Temperature Warnings Reported in an End Cabinet.....	159
About End-Cabinet Cooling.....	160
Ensure End-Cabinet Temperature Set Point is Defined.....	160

About the System Management Workstation (SMW) Software Installation Guide

Release 7.2 UP04

This publication supports Cray software release 7.2 UP04 for Cray System Management Workstations (SMW), released on 24 September 2015.

System Management Workstation (SMW) Software Installation Guide (S-2480-7204) includes the following new content:

- Procedures for installing and configuring the new SMW model: the Dell PowerEdge™ R630 Rack Server
- Troubleshooting procedures for
 - recovering from a damaged or missing MySQL database
 - preventing or correcting inadequate end-cabinet cooling (liquid-cooled systems)

This publication was republished as System Management Workstation (SMW) Software Installation Guide (S-2480-7204a) on 30 October 2015 with the following correction:

- Added instructions to partition the drive and specify 32GB as the swap partition size in [R630 SMW: Install the SLES Base Operating System](#) on page 26 (step 7).

Audience and Scope

This publication is intended for system administrators who are familiar with operating systems derived from Linux.

This publication contains procedures for installing Cray System Management Workstation (SMW) operating software on desk-side and rack-mount (R815 and R630) SMW hardware.

- [Install the SMW Release Package](#) on page 7: How to perform an initial or clean 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.
- [Configure the Boot RAID](#) on page 62: How to configure the system boot RAID from two different vendors: Data Direct Networks, Inc. (DDN) and NetApp, Inc.
- [Update or Upgrade the Cray SMW Software](#) on page 109: How to update or upgrade SMW software.
IMPORTANT: Before proceeding with any upgrade or update installation, read the release notes to determine supported upgrade paths from previous versions of the SMW software.
- [Perform Additional SMW Installation or Configuration Procedures](#) on page 143: How to install third-party software RPM packages, change the time zone, and so forth.
- [Troubleshoot SMW Problems](#) on page 152: How to restore AC power after a power failure, replace a failed disk drive, and others.

Feedback

Your feedback is important to us. Visit the Cray Publications Portal at <http://pubs.cray.com> and make comments online using the **Contact Us** button in the upper-right corner, or email comments to pubs@cray.com.

Related Publications

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

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

Distribution Media

The Cray SMW 7.2UP04 release distribution media consists of three DVDs.

- For an initial installation of the SMW software, use the following DVDs in the order listed:
 1. *Dell Driver Update Disk*, which is used for the R630 SMW only.
 2. *Cray-SMWbase11SP3-201507081500*, which contains the SMW base operating system, SUSE Linux Enterprise Server version 11 Service Pack 3 (SLES 11 SP3). This is used for all SMW models (R630, R815, and desk-side).
 3. *Cray SMW 7.2UP04*, which contains the Cray SMW 7.2 UP04 software packages. This is used for all SMW models.
- For an update of the SMW software, use only the DVD labeled *Cray SMW 7.2UP04*, which contains the Cray SMW 7.2 UP04 software packages. To update to the *Cray 7.2.UP04* release, ensure that the system is running SLES 11 SP3 and a supported update version (see Supported CLE and SMW Upgrade Paths in *Cray SMW 7.2.UP04 Software Release Announcement*).

Install the SMW Release Package

Follow these procedures to perform an initial or clean installation of the Cray System Management Workstation (SMW) 7.2.UP04 release package.

In addition to a desk-side SMW model, Cray provides two rack-mount SMW models: the Dell PowerEdge™ R815 Rack Server and the Dell PowerEdge™ R630 Rack Server. The figure below shows an easy way to distinguish between the two rack-mount models when viewing them from the front.

Figure 1. Distinguishing Features of Dell R815 and R630 Servers



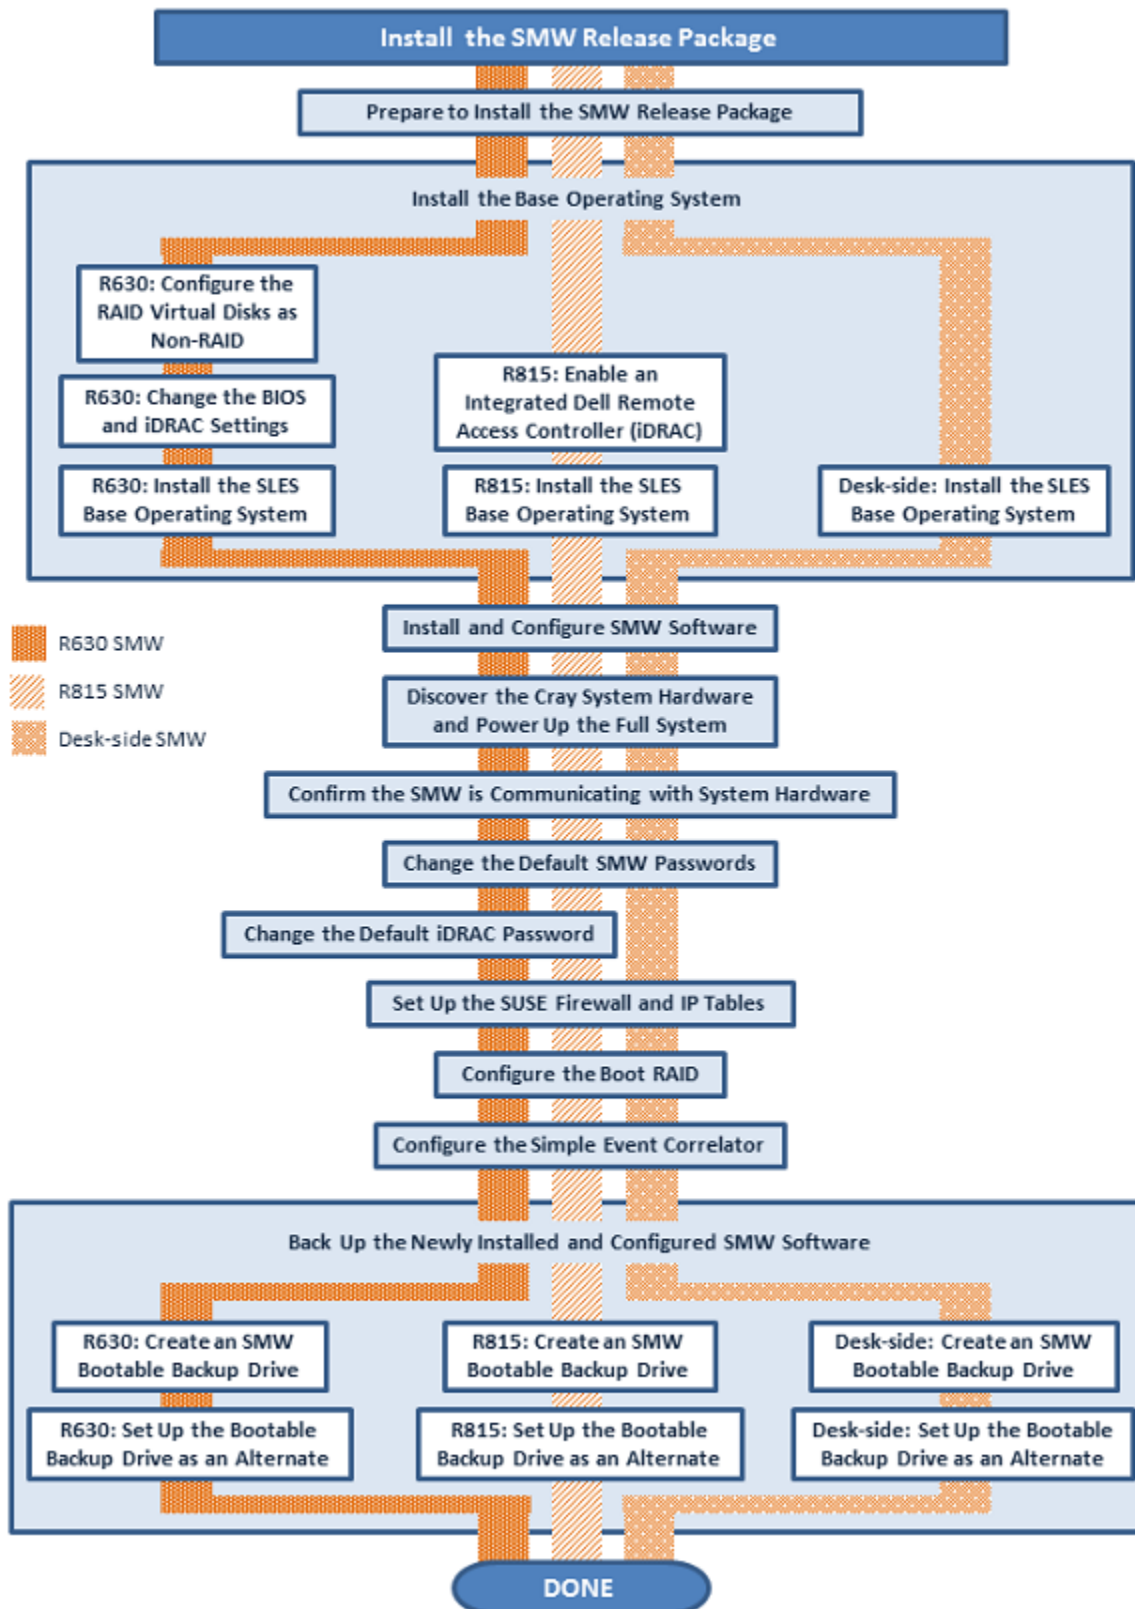
Dell R815: 2U high and 6 drive bays



Dell R630: 1U high and 8 drive bays

Most of the installation procedures that follow apply to all SMW models. Those that apply to only one or more of the models are clearly marked, and the visual guide below also indicates applicability.

Figure 2. Visual Guide to SMW Release Package Installation



Prepare to Install the SMW Release Package

- Read the *SMW Release Errata* and the *SMW README* provided with the 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 (FN) related to kernel security fixes to identify any changes to this release package. Apply any needed changes before installing the new software.
- Use the correct publication. To install the SMW release package on a system configured for SMW high availability (HA) with the SMW failover feature, use *SMW HA XC Installation Guide* (S-0044). Otherwise, use *System Management Workstation (SMW) Software Installation Guide* (S-2480).
- Verify that the network connections are in place (see [Network Connections](#) on page 9).
- Know which configuration values are site-specific and which are defaults (see [Configuration Values](#) on page 9).
- Refer to the default passwords used during the installation process (see [Passwords](#) on page 11).
- Become familiar with the mapping of Linux device names to physical drive slots if using a rack-mount SMW model (see [Rack-mount SMW: Linux /dev Names Mapped to Physical Drive Slots](#) on page 11).

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
 - eth5 - Used for mirrored storage for the power management database (PMDb), if available
 - eth6-7 - Reserved for future use

NOTE: Ethernet port assignments are valid only after the SMW software installation completes (see [R815 SMW: Install the SLES Base Operating System](#) on page 36).

- The SMW must have a Fibre Channel or Serial Attached SCSI (SAS) connection to the boot RAID.
- The boot node must have a Fibre Channel or SAS connection to the boot RAID.
- The service database (SDB) node must have a Fibre Channel or SAS connection to the boot RAID.

IMPORTANT: The SMW must be disconnected from the boot RAID before the initial installation of the SLES software. Ensure that the Fibre Channel optic cable connectors or SAS cable connectors have protective covers.

Configuration Values

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

NOTE: These default IP addresses are only for a standalone SMW. For an SMW HA system, see the default IP addresses in [Configuration Values for an SMW HA System](#).

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, <code>eth1</code>
10.2.1.1	SMW, <code>eth2</code> - Reserved for SMW failover
10.3.1.1	SMW, <code>eth3</code>
10.3.1.254	boot node
10.4.1.1	SMW, <code>eth4</code> - Reserved for SMW failover
127.0.0.1	localhost (loopback)

The following configuration values are site dependent. Record the actual values for the installation site in the third column. References to rack-mount SMW include both the Dell R815 and Dell R630 models..

NOTE: In addition to these values, there are HA-specific values that apply to an SMW HA system. See [Configuration Values for an SMW HA System](#).

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	
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	
For rack-mount SMW only: iDRAC hostname	cray-drac	
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	
Timezone	US/Central	

Description	Example	Actual Value
NTP servers	ntphost1 ntphost2	
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 changing 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; a password must be created
admin (DDN™ boot RAID)	password
user (DDN boot RAID)	password
admin (DDN storage RAID)	password
user (DDN storage RAID)	password
root (iDRAC)	initial0

Rack-mount SMW: Linux /dev Names Mapped to Physical Drive Slots

Due to the nature of the hardware controllers on rack-mount SMWs, the physical drive slots do not have a fixed logical mapping to the Linux `/dev/sd?` device names. In the sections that follow, tables show the mapping of Linux `/dev` names map to physical drive slots and the physical placement of those drives for the current rack-mount SMW models.

The information contained in these tables relates to the `SMWinstall.conf` file and is useful to know before performing a system backup.

R630 SMW: Mapping of Linux /dev Names to Physical Drive Slots

IMPORTANT: Due to the nature of the hardware controllers on the R630 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 for the R630 SMW

Drive Slot	Linux <code>/dev</code> Name
0	<code>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0</code>
1	<code>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0</code>
2	<code>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:2:0</code>
3	<code>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:3:0</code>
4	<code>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:4:0</code>
5	<code>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:5:0</code>
6	<code>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:6:0:0</code>
7	<code>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:7:0:0</code>

Table 5. Physical Placement of R630 Rack-mount SMW Chassis Hard Drives (front view)

		Slot 2	Slot 4	Slot 6
Slot 0	Slot 1	Slot 3	Slot 5	Slot 7

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: empty
- Slot 6: empty
- Slot 7: empty

R815 SMW: Mapping of Linux /dev Names to Physical Drive Slots

IMPORTANT: Due to the nature of the hardware controllers on the R815 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 6. Mapping of Linux `/dev` Names to Physical Drive Slots for the R815 SMW

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 7. Physical Placement of R815 Rack-mount SMW Chassis Hard Drives (front view)

Slot 0 phy7 SLES/SMW SW (160 GB)	Slot 2 phy5 Database (160 GB)	Slot 4 phy3 optional drive for PMDB
Slot 1 phy6 SLES/SMW SW (160 GB)	Slot 3 phy4 log/debug/dump (500 GB or 1 TB)	Slot 5 phy2 Empty

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: Empty

About the System Management Workstation (SMW) Base Operating System Installation Process

The base operating system must be installed on the SMW before Cray SMW software is installed.

Installation of the Operating System on a Dell R630 SMW

To install the base operating system on a Dell R630 SMW, use the following procedures in the order listed:

1. Configure the SMW disks as non-RAID. See [R630 SMW: Configure the RAID Virtual Disks as Non-RAID](#) on page 14 for instructions.

2. Configure the BIOS and iDRAC for the SMW. See [R630 SMW: Change the BIOS and iDRAC Settings](#) on page 16 for instructions.
3. Install the base operating system. See [R630 SMW: Install the SLES Base Operating System](#) on page 26 for instructions.

Installation of the Operating System on a Dell R815 SMW

To install the base operating system on a Dell R815 SMW, use the following procedures in the order listed:

1. Configure the BIOS and iDRAC and enable iDRAC for the SMW. See [R815 SMW: Enable an Integrated Dell Remote Access Controller \(iDRAC\)](#) on page 31 for instructions.
2. Install the base operating system. See [R815 SMW: Install the SLES Base Operating System](#) on page 36 for instructions.

Installation of the Operating System on a Desk-side SMW

To install the base operating system on a desk-side SMW, use the following procedure: [Desk-side SMW: Install the SLES Base Operating System](#) on page 41

R630 SMW: Configure the RAID Virtual Disks as Non-RAID

About this task

This procedure configures the RAID virtual disks so that the internal Dell PERC RAID controller treats these disks as non-RAID. Converting RAID-capable disks to non-RAID is necessary for R630 SMWs with SLES11SP3 as the base operating system. Because Cray ships systems with most of the installation and configuration completed, some of the steps are optional and are needed only to change the configuration.

This procedure assumes a Dell R630 server and the PERC H330 Mini BIOS Configuration Utility 4.03-0010. For a different server model and version of RAID configuration utility, minor adjustments to these steps may be necessary. For more information, refer to the documentation for the relevant Dell PERC controller or server RAID controller software.

Procedure

1. Connect a keyboard, monitor, and mouse to the USB and monitor connectors on the SMW, if not already connected.
2. Ensure all SMW internal disk drives are inserted into the SMW drive slots.
3. Power up the SMW, and when prompted, press **Ctrl-R** to enter the RAID controller configuration utility.

Cray recommends using the RAID configuration utility (via **Ctrl-R**) instead of the **System Setup Device Settings** menu to configure the RAID virtual disks.

TIP: In the RAID configuration utility:

- Use the up-arrow or down-arrow key to highlight an item in a list, then press the **Enter** key to select the item.
- Press the **F2** key to display a menu of options for an item.
- Use the right-arrow, left-arrow, or **Tab** key to switch between the **Yes** and **No** buttons in a confirmation window.

4. Clear existing/default disk configuration, if necessary.
 - a. Select `Disk Group 0` and press **F2** if any disk groups are currently defined.
 - b. Select `Delete Disk Group` and then press **Enter**.
 - c. Select `Yes` in the pop-up confirmation window to confirm the changes.
5. Configure all disks as non-RAID.
 - a. Select `No Configuration Present` and then press **F2**.
 - b. Select `Convert to Non-RAID` and then press **Enter**.

The `Convert RAID Capable Disks to Non-RAID` screen opens.

NOTE: If all of the drives are already non-RAID, the menu items will appear unavailable. In this case, proceed to step 6 on page 15.

- c. Select all physical disks for this RAID-5 disk group: press **Enter** to select the checkbox for a physical disk and advance to the next disk, and repeat until each physical disk is selected.
 - d. Press **Tab** to move to `OK`, then press **Enter**.
6. Verify the virtual disk changes.

Compare settings with those shown in the figure to verify the virtual disk changes (`Disk Groups: 0`, `Virtual Disks: 0`, `Physical Disks: 5`).

Figure 3. Final RAID Configuration Settings



7. Switch disk controller from RAID-Mode to HBA-mode.
 - a. Press **Ctrl-N** to move from the `VD Mgmt` tab to the `PD Mgmt` tab and then to the `Ctrl Mgmt` tab.

- b. Press **Tab** to move to `Personality Mode`, then press **Enter**.

A choice between `RAID-Mode` and `HBA-mode` appears.

- c. Use the down-arrow key to select `HBA-Mode`, then press **Enter**.

- d. Press **Tab** to move to `Apply`, then press **Enter**.

This message appears:

```
The operation has been performed successfully. Reboot the system for the
change to take effect.
```

8. Exit RAID configuration utility.

- a. Press the **Escape** key to exit the RAID configuration utility.

- b. Press **Tab** to move to `OK`, then press **Enter** to confirm.

A message appears prompting a reboot. This action to reboot occurs as the first step of the next procedure.

```
** Press Control+Alt+Delete to reboot **
```

R630 SMW: Change the BIOS and iDRAC Settings

About this task

This procedure describes how to change the system setup for BIOS and embedded server (iDRAC) settings on the R630 SMW node. Because Cray ships systems with most of the installation and configuration completed, some of the steps are optional and are needed only to change the configuration.

This procedure assumes a Dell R630 server and the PERC H330 Mini BIOS Configuration Utility 4.03-0010. For a Dell R815 server, see [R815 SMW: Change the BIOS and iDRAC Settings for an iDRAC](#) on page 31.

Procedure

1. Press **Ctrl-Alt-Delete** to reboot.


The server will restart the boot process; however that will not interrupt RAID initialization. Watch as the system reboots and the BIOS power-on self-test (POST) process begins. Be prepared to press **F2**, when prompted, to change the system setup.

2. Press the **F2** key immediately 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 color from white-on-black to white-on-blue.

After the POST process completes and all disk and network controllers have been initialized, the Dell System Setup screen appears. The following submenus are available on the System Setup Main Menu and will be used in subsequent procedures: `System BIOS`, `iDRAC Settings`, and `Device Settings`.

Figure 4. Dell R630 System Setup Main Menu

TIP: In system setup screens,

- Use the **Tab** key to move to different areas on the screen.
- Use the up-arrow and down-arrow keys to highlight or select an item in a list, then press the **Enter** key to enter or apply the item.
- Press the **Esc** key to exit a submenu and return to the previous screen.

3. Change the BIOS settings.

- a. Select **System BIOS** on the **System Setup Main Menu**, then press **Enter**.

The **System BIOS Settings** screen appears.

Figure 5. Dell R630 System BIOS Settings Screen

b. Change Boot Settings.

1. Select **Boot Settings** on the System BIOS Settings screen, then press **Enter**. The Boot Settings screen appears.

Figure 6. Dell R630 Boot Settings Screen

2. Ensure that **Boot Mode** is BIOS and not UEFI.
3. Select **Boot Option Settings**, then press **Enter**.

4. Select **Boot Sequence** on the **Boot Option Settings** screen, then press **Enter** to view a pop-up window with the boot sequence.


Figure 7. Dell R630 BIOS Boot Sequence



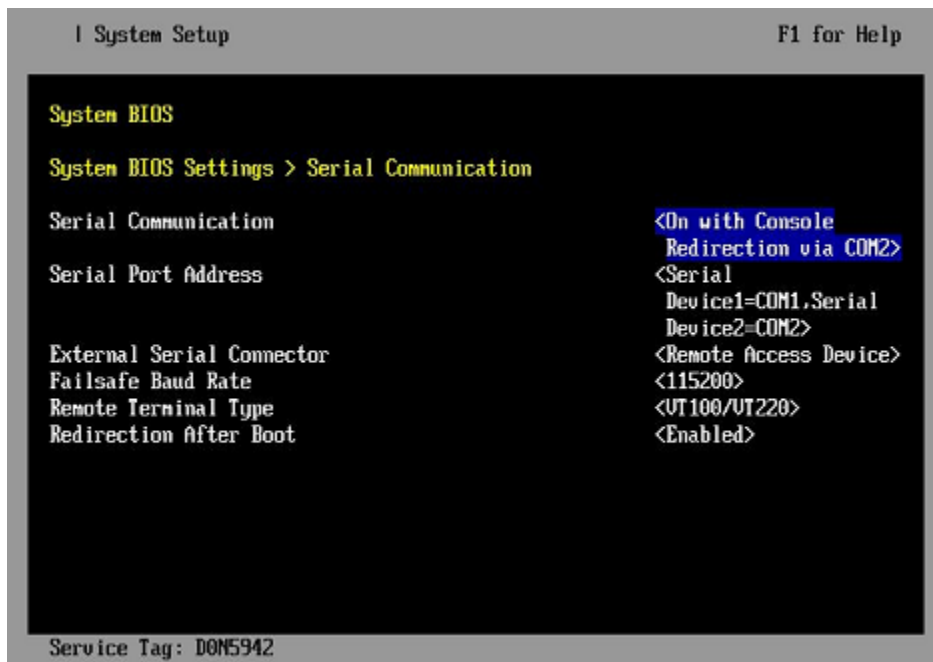
5. Change the boot order in the pop-up window so that the optical drive appears first, then the hard drive. If **Integrated NIC** appears in the list, it should end up below the optical drive and hard drive in the list.

TIP: Use the up-arrow or down-arrow key to highlight or select an item, then use the + and - keys to move the item up or down.

6. Select **OK**, then press **Enter** to accept the change.
 7. Click the box next to **Hard drive C:** under the **Boot Option/Enable/Disable** section to enable it. Do the same for the optical drive, if necessary.
 8. Select **integrated NIC**, then press **Enter** to disable it.
 9. Press **Esc** to exit **Boot Option Settings**.
 10. Press **Esc** to exit **Boot Settings** and return to the **System BIOS Settings** screen.
- c. Change Serial Communication Settings.

Figure 8. Dell R630 System BIOS Settings: Serial Communication

1. Select **Serial Communication** on the **System BIOS Settings** screen. The **Serial Communication** screen appears.


Figure 9. Dell R630 Serial Communication Screen

2. Select **Serial Communication** on the **Serial Communication** screen, then press **Enter**. A pop-up window displays the available options.
3. Select **On with Console Redirection via COM2** in the pop-up window, then press **Enter** to accept the change.

4. Select `Serial Port Address`, then select `Serial Device1=COM1`, `Serial Device2=COM2`, then press **Enter**.
 5. Select `External Serial Connector`, then press **Enter**. A pop-up window displays the available options.
 6. Select `Remote Access Device` in the pop-up window, then press **Enter** to return to the previous screen.
 7. Select `Failsafe Baud Rate`, then press **Enter**. A pop-up window displays the available options.
 8. Select `115200` in the pop-up window, then press **Enter** to return to the previous screen.
 9. Press the **Esc** key to exit the `Serial Communication` screen.
 10. Press the **Esc** key to exit the `System BIOS Settings` screen. A "Settings have changed" message appears.
 11. Select `Yes` to save changes. A "Settings saved successfully" message appears.
 12. Select `Ok`.
4. Change the iDRAC (Integrated Dell Remote Access Controller) settings.
- a. Select `iDRAC Settings` on the `System Setup Main Menu`, then press **Enter**.

The iDRAC Settings screen appears.

Figure 10. Dell R630 iDRAC Settings Screen




- b. Change the network.
 1. Select `Network` to display a long list of network settings.
 2. Change the DNS DRAC name: use the down-arrow key to scroll to `DNS DRAC Name`, then enter an iDRAC host name that is similar to the SMW node host name (e.g., `cray-drac`).
 3. Change the static DNS domain name: use the down-arrow key to scroll to `Static DNS Domain Name`, enter the DNS domain name, then press **Enter**.

4. Change the IPv4 Settings.
 - a. Use the down-arrow key to scroll to the `IPV4 SETTINGS` list.
 - b. Ensure that IPv4 is enabled.
 - i. If necessary, select `Enable IPV4` and press **Enter**.
 - ii. Select `<Enabled>` in the pop-up window.
 - iii. Press **Enter** to return to the previous screen.
 - c. Ensure that DHCP is disabled.
 - i. If necessary, select `Enable DHCP` and press **Enter**.
 - ii. Select `<Disabled>` in the pop-up window.
 - iii. Press **Enter** to return to the previous screen.
 - d. Change the IP address.
 - i. Select `Static IP Address`.
 - ii. Enter the IP address of the iDRAC interface (`ipmi0`) for the SMW and press **Enter**.
 - e. Change the gateway.
 - i. Select `Static Gateway`.
 - ii. Enter the appropriate value for the gateway of the network to which the iDRAC is connected and press **Enter**.
 - f. Change the subnet mask.
 - i. Select `Subnet Mask`.
 - ii. Enter the subnet mask for the network to which the iDRAC is connected (such as `255.255.255.0`) and press **Enter**.
 - g. Change the DNS server settings.
 - i. Select `Static Preferred DNS Server`.
 - ii. Enter the IP address of the primary DNS server and press **Enter**.
 - iii. Select `Alternate DNS Server`.
 - iv. Enter the IP address of the alternate DNS server and press **Enter**.
 5. Change the IPMI settings to enable the Serial Over LAN (SOL) console.
 - a. Use the down-arrow key to scroll to the `IPMI SETTINGS` list.
 - b. If necessary, select `Enable IPMI over LAN` to ensure that IPMI over LAN is enabled.

TIP: Use the left-arrow or right-arrow to switch between two settings.
 - c. If necessary, select `Channel Privilege Level Limit to be Administrator` to ensure that Channel Privilege Level Limit is set to Administrator.
 6. Exit Network screen: Press the **Esc** key to exit the `Network` screen and return to the `iDRAC Settings` screen.
- c. Change Front Panel Security to show the host name in LCD display.
 1. Use the down-arrow key to scroll down and highlight `Front Panel Security` on the `iDRAC Settings` screen, then press **Enter**.
 2. Select `Set LCD message`, then press **Enter**.

3. Select `User-Defined String`, then press **Enter**.
 4. Select `User-Defined String`, then enter the SMW host name and press **Enter**.
 5. Press the **Esc** key to exit the `Front Panel Security` screen.
- d. (Optional) Change the `System Location` configuration on the `iDRAC Settings` screen to set any of these fields: `Data Center Name`, `Aisle Name`, `Rack Name`, and `Rack Slot`.
- e. Change `User Configuration`.
1. Use the down-arrow key to highlight `User Configuration` on the `iDRAC Settings` screen, then press **Enter**.
 2. Confirm that `User Name` is `root`: select `User Name`, then enter the "root" user name.
 3. Set a password.
 - a. Select `Change Password`, then enter a new password.
 - b. Reenter the new password in the next pop-up window to confirm it (the default password is "calvin").
 4. Press the **Esc** key to exit the `User Configuration` screen.
- f. Exit `iDRAC settings`.
1. Press the **Esc** key to exit the `iDRAC Settings` screen. A "Settings have changed" message appears.
 2. Select `Yes`, then press **Enter** to save the changes. A "Success" message appears.
 3. Select `Ok`, then press **Enter**. The main screen (`System Setup Main Menu`) appears.
5. Change device settings.
- These steps disable an integrated NIC device by changing the setting for the integrated NIC on a port from `PXE` to `None`.
- a. Change `Integrated NIC 1 Port 1`
1. Select `Device Settings` on the `System Setup Main Menu`, then press **Enter**. The `Device Settings` screen appears.

Figure 11. Dell R630 Device Settings Screen



2. Select Integrated NIC 1 Port 1: ... on the Device Settings screen, then press **Enter**.
3. Select MBA Configuration Menu on the Main Configuration Page screen, then press **Enter**.

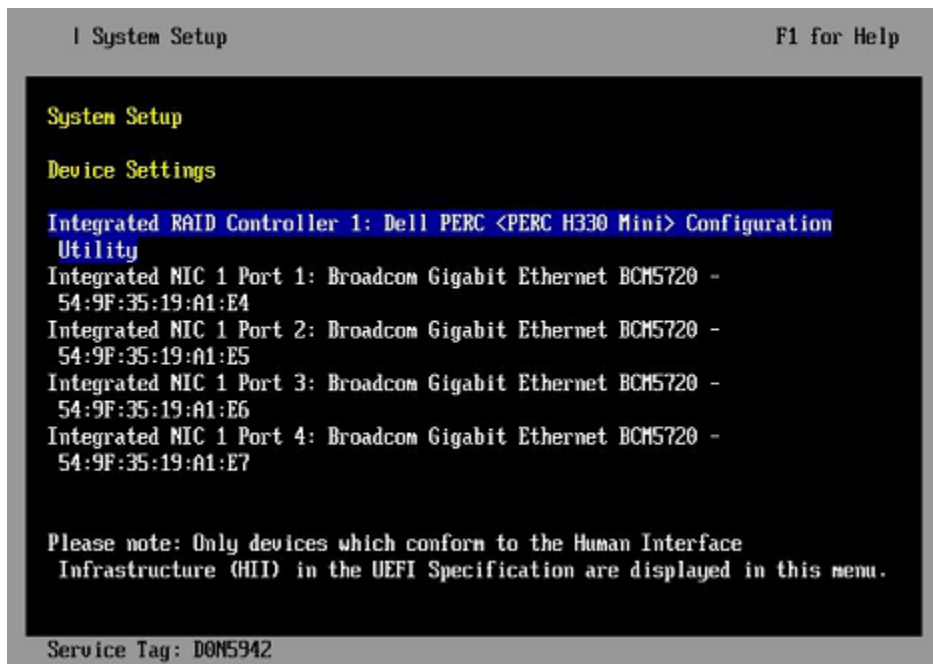
Figure 12. Dell R630 BIOS MBA Configuration Settings



4. Select Legacy Boot Protocol on the MBA Configuration Menu screen, use the right-arrow or left-arrow key to highlight None, and press **Enter**.
5. Press the **Esc** key to exit the MBA Configuration Menu screen.

6. Press the **Esc** key to exit the Main Configuration Page screen. A "Warning Saving Changes" message appears.
 7. Select **Yes**, then press **Enter** to save the changes. A "Success" message appears.
 8. Select **OK**, then press **Enter**. The Device Settings screen appears.
 9. Press the **Esc** key to exit the Device Settings screen. A "Settings have changed" message appears.
 10. Select **Yes**, then press **Enter** to save the changes. A "Settings saved successfully" message appears.
 11. Select **OK**, then press **Enter**. The main screen (System Setup Main Menu) appears.
- b. Change Integrated NIC 1 Port 2
1. Select Device Settings on the System Setup Main Menu, then press **Enter**. The Device Settings screen appears.

Figure 13. Dell R630 Device Settings Screen



2. Select Integrated NIC 1 Port 2: ... on the Device Settings screen, then press **Enter**.
3. Select MBA Configuration Menu on the Main Configuration Page screen, then press **Enter**.

Figure 14. Dell R630 BIOS MBA Configuration Settings



4. Select **Legacy Boot Protocol** on the MBA Configuration Menu screen, use the right-arrow or left-arrow key to highlight **None**, and press **Enter**.
5. Press the **Esc** key to exit the MBA Configuration Menu screen.
6. Press the **Esc** key to exit the Main Configuration Page screen. A "Warning Saving Changes" message appears.
7. Select **Yes**, then press **Enter** to save the changes. A "Success" message appears.
8. Select **OK**, then press **Enter**. The Device Settings screen appears.
9. Press the **Esc** key to exit the Device Settings screen. A "Settings have changed" message appears.
10. Select **Yes**, then press **Enter** to save the changes. A "Settings saved successfully" message appears.
11. Select **OK**, then press **Enter**. The main screen (System Setup Main Menu) appears.

R630 SMW: Install the SLES Base Operating System

About this task

This procedure describes how to install SLES11SP3 as the base operating system on an R630 SMW node. For a Dell R815 server, see [R815 SMW: Install the SLES Base Operating System](#) on page 36.

Procedure

1. Remove all SMW internal disk drives except drive 0, and disconnect the boot RAID from the SMW.
This must be done before running the installer.
2. Insert the Dell Driver Update Disk (DUD) DVD in the optical (DVD) drive of the SMW.

(The DVD drive on the front of the SMW may be hidden by a removable decorative bezel.)

3. Exit Dell System Setup.

- a. Press the **Esc** key while viewing the System Setup Main Menu to exit the Dell System Setup utility.

A message appears asking whether to exit and reboot.

- b. Select **Yes**, then press **Enter**.

The server will restart the boot process.

4. Boot from the Dell DUD DVD (dell_shared_perc_8-sle11sp3-x86_64-1.0 DVD).

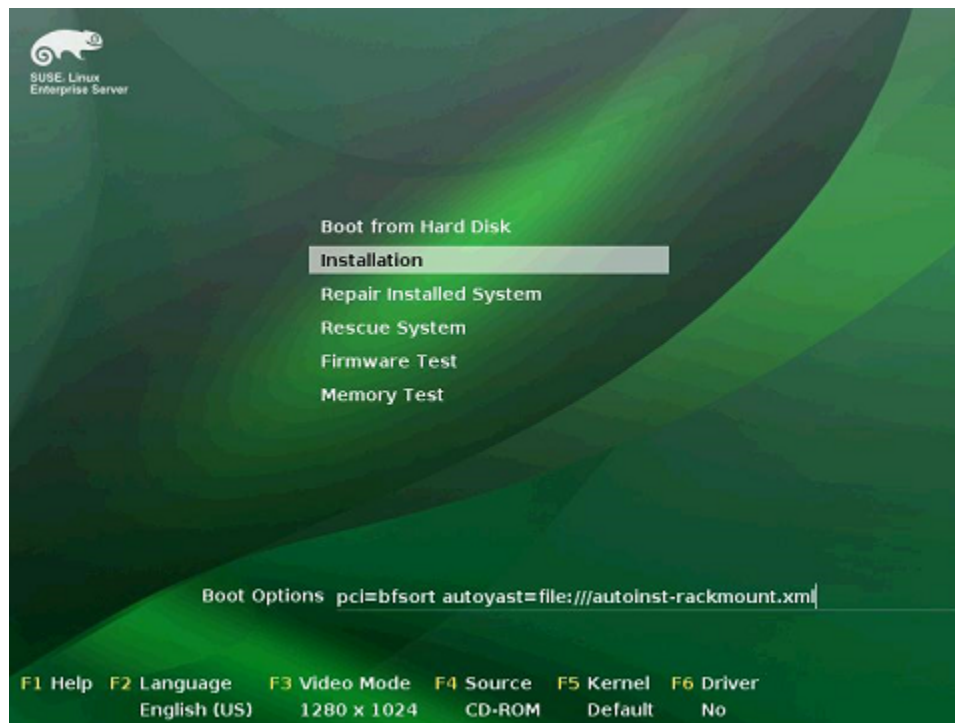
5. Select installation.

- a. Select **Installation** rather than the default action of booting from the hard disk.

```
- Boot from Hard Disk
- Installation
- Repair Installed System
- Rescue System
- Firmware Test
- Memory Test
```

There is a brief timeout of 10 to 15 seconds before the default action is taken. Continue to the next step quickly before the timeout ends. If the timeout ends before a selection is made, the system will boot from the hard disk (the default selection). If that happens, shut down the SMW, then begin the power-up sequence again.

Figure 15. Dell DUD Installation Screen: Installation Boot Options



- b. Enter the following line in the `Boot Options` text input box, and then press **Enter** to begin the installation.

```
pci=bfsort autoyast=file:///autoinst-rackmount.xml
```

A status window will appear showing that the Linux Kernel is booting.

```
Starting...
```

```
Loading Linux Kernel
```


6. Load the installation.

- a. Prepare installation disc.

On the `DELL SHARED PERC 8 Bootable Driver Kit` screen, the following message appears:

```
Please insert disc 1 of the installation media for SUSE Linux Enterprise
Server/Desktop 11 SP3
```

Figure 16. Dell Shared PERC 8 Driver Kit Screen



If `Cray-SMWbase11SP3-201507081500.iso` has not already been burned to a DVD, do that now.

- b. Insert the `Cray-SMWbase11SP3-201507081500` DVD into the DVD drive, then press **Enter**.

```
Loading Installation System (1/6)
Loading Installation System (2/6)
Loading Installation System (3/6)
Loading Installation System (4/6)
Loading Installation System (5/6)
Loading Installation System (6/6)
```

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

- a. Remove the `sda` partitions.

1. For each partition of `sda`, double-click the `Device name` of that specific `Hard Disk` table entry. This opens the `Partition: /dev/sd...` table for that drive, with `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 `sda` disk until the displayed list is empty.

- b. Create new partitions for `swap` and `root` for the `sda` device sized for this 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. The `Add Partition on /dev/sd...` screen appears.
 2. Select `Primary Partition`, then select `Next`. The `New Partition Size` screen appears.
 3. Select `Custom size`, enter **32GB** as the swap partition size, then select `Next`. The `Formatting Options and Mounting Options` screen appears.
 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 appears again, but with the new partition entry.
 7. Select the **Add** button at the bottom of the window. The `Add Partition on /dev/sd...` screen appears.
 8. Select `Primary Partition`, then select `Next`. The `New Partition Size` screen appears.
 9. Select `Maximum Size`, then select `Next`. The `Formatting Options and Mounting Options` screen appears.
 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 is complete. On the `Expert Partitioner` screen, select `Accept` to accept the changes. A pop-window stating the following appears:

```
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.

- c. Confirm the disk partitions.

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

```
Partitioning
Create root partition /dev/sda2 (146.92 GB) with ext3
use /dev/sda1 as swap
```

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

8. Run the installation.

After the install program does some system probing and setup for the install, the `Installation Settings` window appears.

- a. Select **Install**.

The license agreement window appears.

- b. Select the **I Agree** button in the license agreement window to continue the installation process.

```
YaST2
Confirm Package License: agfa-fonts
(long detailed license agreement)
```

- c. Select **Install** in the **Confirm Installation** window 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.

9. When the **Network Settings** window for configuring the customer network appears, select the entry labeled **eth0 Customer Network Ethernet**. On the **Overview** 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. Enter the hostname and domain name separately.

Host Name	Domain Name
smwhost	my.domain.com

If a fully qualified hostname that includes the domain name is entered, the hostname is accepted but the periods are removed; 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**.

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

- a. Select the appropriate 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.

11. Remove the protective covers from the Fibre Channel (FC) or SAS cables and connectors, clean the ends of the cables and connectors, and reconnect the data cables.

12. Re-seat the previously ejected SMW internal disk drives.

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

```
smw# reboot
```

Installation of the base operating system for the R630 SMW is now complete. The system is now ready for installation and configuration of the SMW software. Go to [Install and Configure SMW Software](#) on page 45.

R815 SMW: Enable an Integrated Dell Remote Access Controller (iDRAC)

Enabling an Integrated Dell Remote Access Controller (iDRAC) allows an administrator to manage a rack-mount SMW remotely. For a Dell R815 server, it is a two-step process:

1. Change the BIOS and iDRAC settings. See [R815 SMW: Change the BIOS and iDRAC Settings for an iDRAC](#) on page 31 for instructions.
2. Enable the iDRAC. See [R815 SMW: Enable an iDRAC](#) on page 35 for instructions.

Before enabling an iDRAC on an R815 SMW, the system administrator must:

- Have physical access to the SMW console.
- Know the iDRAC IP address, subnet mask, and default gateway.
- Know the SMW `root` account password.

R815 SMW: Change the BIOS and iDRAC Settings for an iDRAC

About this task

Follow this procedure to change the BIOS and iDRAC settings for a Dell R815 SMW.

NOTE: The output and figures shown in this procedure are used as examples. Actual output and display may vary.

Procedure


1. 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`.

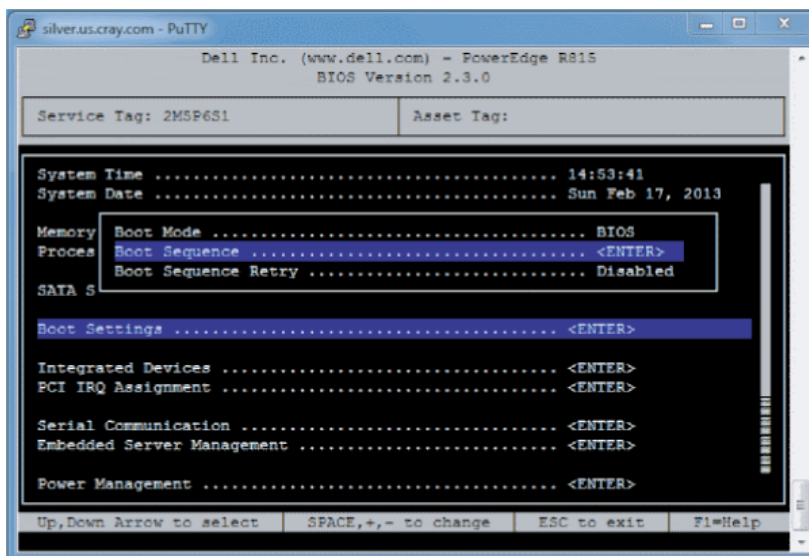
2. Select **Boot Settings**, then press `Enter`.

Figure 17. Dell R815 SMW Boot Settings Menu



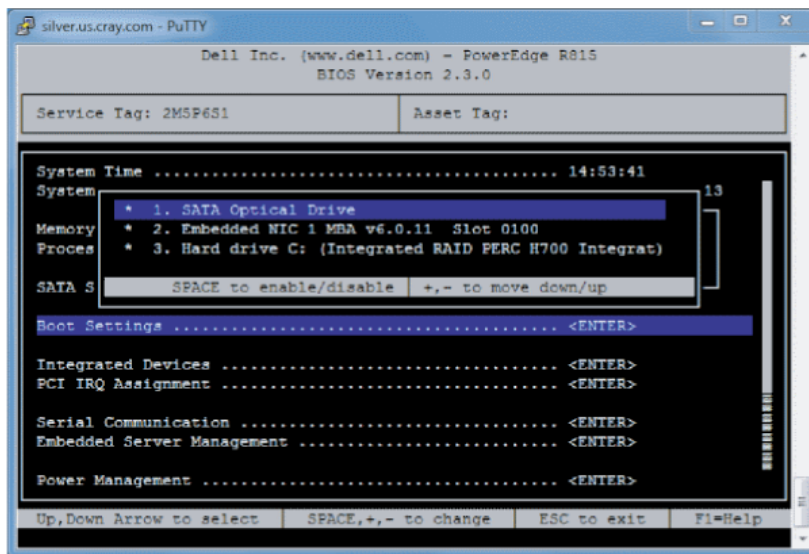
- a. Select **Boot Sequence**, then press **Enter** to view the boot settings.

Figure 18. Dell R815 SMW Boot Sequence Menu



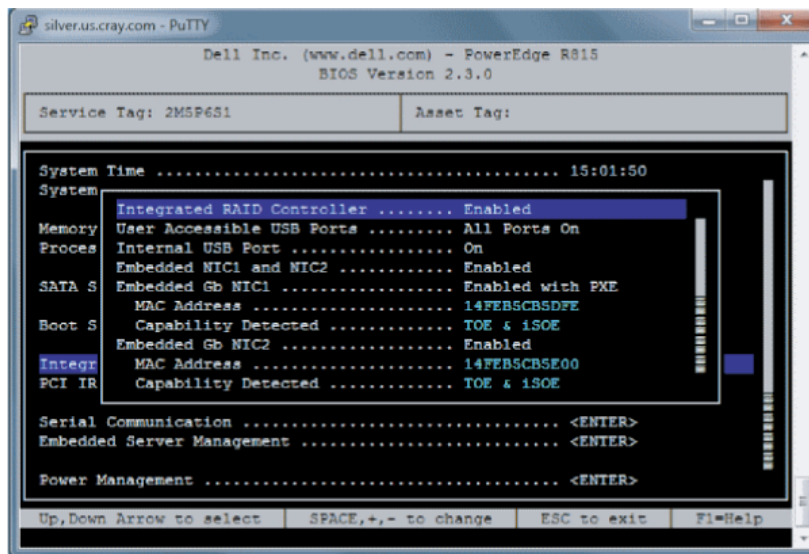
- b. In the pop-up window, change the boot order so that the optical (DVD) drive appears first, then the hard drive. If integrated NIC appears in the list, it should end up below the optical drive and hard drive in the list.

Figure 19. Dell R815 SMW Boot Sequence Settings




- c. Select integrated NIC, then press **Enter** to disable it.
 - d. Press **Enter** to return to the **BIOS Boot Settings** screen.
3. Press **Esc** to return to the System Setup Menu, scroll down and select **Integrated Devices**.

Figure 20. Dell R815 SMW Integrated Devices (NIC) Settings



- a. Set **Embedded Gb NIC 2** to **Enabled**.
 - b. Scroll down and set **Embedded NIC 3** to **Enabled**.
 - c. Set **Embedded Gb NIC 4** to **Enabled**.
 - d. Press **Esc** to return to the System Settings Menu.
4. Select **Embedded Server Management**.

Figure 21. Dell R815 SMW Embedded Server Management Settings




- a. Set **Front-Panel LCD Options** to **User-Defined LCD String**.
- b. Set **User-Defined LCD String** to the login hostname, such as `cray-drac`.
5. Save the changes and exit.
 - a. Press `Esc` to exit the System Setup Main Menu.
 - b. The utility displays the prompt "Are you sure you want to exit and reboot?" Select **Yes**.
6. When the system reboots, press `Ctrl-E` to configure the iDRAC port settings.

www.dell.com

```
iDRAC6 Configuration Utility 1.60
Copyright 2011 Dell Inc. All Rights Reserved
Four 2.10 GHz Twelve-core Processors, L2/L3 Cache: 6 MB/10 MB
iDRAC6 FirmwareRevision: 1.70.21
.
.
.
IPv4 Stack      : Enabled
IP Address      : 10.148. 0 . 2
Subnet mask     : 255.255. 0 . 0
Default Gateway : 0 . 0 . 0 . 0
Press <Ctrl-E> for Remote Access Setup within 5 sec.....
```

- a. Set the **iDRAC LAN** to **ON**.
- b. Select **LAN Parameters** and press `Enter`. Set the IPv4 address to the SMW DRAC IP address.
- c. Press `Esc` to return to the iDRAC menu, and `Esc` to exit and save.

Figure 22. Dell R815 SMW DRAC IPv4 Parameter Settings



R815 SMW: Enable an iDRAC

Procedure

1. Connect Ethernet cable to the iDRAC port. The cable is located on back of the R815 SMW in the lower left corner.
2. Power up the SMW.
3. After the BIOS and disk map have displayed, the IPv4/IPv6 information displays. When the IPv4/IPv6 information displays, press **Ctrl-E**.

NOTICE: The remaining steps in this procedure are similar to step 6 of the previous procedure, [R815 SMW: Change the BIOS and iDRAC Settings for an iDRAC](#) on page 31, which also instructs the user to "press **Ctrl-E**" and configure iDRAC port settings. Review the following steps to see if they have already been done.

4. Using the arrow keys, select **LAN Parameters**, then press **Enter**.
5. Select **NIC Selection** and set it to **Dedicated**. Then press **Esc**.
6. 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 the iDRAC 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.
7. 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: *****
```
- b. Press `Esc`.
8. Press `Esc` again.
9. Select **Save Changes and Exit**, then press `Enter`. The SMW will complete booting up; no user interaction is required.

## R815 SMW: Install the SLES Base Operating System

### About this task

This procedure describes how to install SLES11SP3 as the base operating system on an R815 SMW. Use the DVD labeled `Cray-SMWbase11SP3-` to perform this installation.


### Procedure

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 Fibre Channel or SAS cables and connectors (if present).

*Figure 23. Dell R815 SMW Rear Connections*



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 setup 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-SMWbase11SP3-201507081500** 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 setup.

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 iDRAC6 Configuration Utility 1.57 appears. 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, then press the **Enter** key.

```
- 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
```

If the timeout ends before a selection is made, the system will boot from the hard disk (the default selection). If that happens, shut down the SMW, then begin the power-up sequence again.

As the base installation progresses, the following phrases appear on the 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. Remove the `sda` partitions.
    1. For each partition of `sda`, double-click the Device name of that specific `Hard Disk` table entry. This opens the `Partition: /dev/sd...` table for that drive, with `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 `sda` disk until the displayed list is empty.
  - b. Create new partitions for `swap` and `root` for the `sda` device sized for this 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. The `Add Partition on /dev/sd...` screen appears.
    2. Select `Primary Partition`, then select `Next`. The `New Partition Size` screen appears.
    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 the SMW.

**TIP:** 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 appears.

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 appears again, but with the new partition entry.
7. Select the `Add` button at the bottom of the window. The `Add Partition on /dev/sd...` screen appears.
8. Select `Primary Partition`, then select `Next`. The `New Partition Size` screen appears.
9. Select `Maximum Size`, then select `Next`. The `Formatting Options and Mounting Options` screen appears.
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 is complete. On the `Expert Partitioner` screen, select `Accept` to accept the changes. A pop-window stating the following appears:

```
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.

c. Confirm the disk partitions.

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

```
Partitioning
Create root partition /dev/sda2 (146.92 GB) with ext3
use /dev/sda1 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 a language from the drop-down menu. Select multiple secondary languages, if desired. 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. 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. When the `Network Settings` window for configuring the customer network appears, select the entry labeled `eth0 Customer Network Ethernet`. On the `Overview` 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. Enter the hostname and domain name separately.

| Host Name | Domain Name   |
|-----------|---------------|
| smwhost   | my.domain.com |

If a fully qualified hostname that includes the domain name is entered, the hostname is accepted but the periods are removed; 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 the appropriate 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.

**27.** Remove the protective covers from the Fibre Channel (FC) or SAS cables and connectors, clean the ends of the cables and connectors, and reconnect the data cables.

**28.** Re-seat the previously ejected SMW internal disk drives.

**29.** Reboot the SMW to allow the SMW to discover the drives properly.

```
smw# reboot
```

Installation of the base operating system for the R815 SMW is now complete. The system is now ready for installation and configuration of the SMW software. Go to [Install and Configure SMW Software](#) on page 45.

## Desk-side SMW: Install the SLES Base Operating System

### About this task

This procedure describes how to install `SLES11SP3` as the base operating system on a desk-side SMW. Use the DVD labeled `Cray-SMWbase11SP3-201507081500` to perform this installation.


### Procedure

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).

Figure 24. Desk-side SMW Rear Connections



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 this command on any Linux system to determine the current UTC:

```
linux> date -u
```

10. Insert the base operating system DVD labeled `Cray-SMWbase11SP3-201507081500` 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 the 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 the 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`, 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 a language from the drop-down menu. Select multiple secondary languages, if desired. 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. 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. Enter the hostname and domain name separately. If a fully-qualified hostname that includes the domain name is entered, 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 the appropriate 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.

**20.** 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.

**21.** Reboot the SMW to allow the SMW to discover the drives properly.

```
smw# reboot
```

## Install and Configure SMW Software

### About this task

This procedure takes approximately 30 minutes.

### Procedure

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-SMWbase11SP3-201507081500`) is still in the CD/DVD drive, eject it.

```
smw# eject
```

3. Place the `Cray 7.2UP04 Software` DVD in the drive and mount it.

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

**NOTE:** If problems occur while 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 the installation site.

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

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

- a. Specify the system interconnection network type. There is no default setting.
- b. Adjust LLM settings, as needed.

By default, the `SMWinstall` program enables LLM. To send the logs from the SMW to a site loghost, adjust these settings: `LLM_siteloghost`, `LLM_sitecompatmode`, and `LLM_altrelay`. If the boot RAID controller IP addresses do not start the specified pattern, adjust `llm_raid_ip`.

- c. Leave the Linux `sar` service on the SMW enabled (default), or disable it, as needed.
- d. Leave the controller log forwarding on the SMW enabled (default), or disable it, as needed.
- e. For rack-mount SMWs only: Set the `LOGDISK`, `DBDISK`, and `PMDISK` variables.

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. For the appropriate disk name for rack-mount models, see the `LOGDISK` section of the `SMWinstall.conf` file shown below.
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. For the appropriate disk name for rack-mount models, see the `DBDISK` section of the `SMWinstall.conf` file shown below.
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. For the appropriate disk name for rack-mount models, see the `PMDISK` section of the `SMWinstall.conf` file shown below.

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 relevant lines in the `SMWinstall.conf` file:

```
#For SLES 11 SP3-based rackmount R815 SMWs:
#LOGDISK=/dev/disk/by-path/pci-0000:05:00.0-sas-phy4-0x4433221104000000-
lun-0
#For SLES 11 SP3-based rackmount R630 SMWs:
#LOGDISK=/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:3:0

#LOGMOUNT=/var/opt/cray/disk/1

#For SLES 11 SP3-based rackmount R815 SMWs:
#DBDISK=/dev/disk/by-path/pci-0000:05:00.0-sas-phy5-0x4433221105000000-
lun-0
#For SLES 11 SP3-based rackmount R630 SMWs:
#DBDISK=/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:2:0

#DBMOUNT=/var/lib/mysql

#For SLES 11 SP3-based rackmount R815 SMWs:
#PMDISK=/dev/disk/by-path/pci-0000:05:00.0-sas-phy3-0x4433221103000000-
lun-0
#For SLES 11 SP3-based rackmount R630 SMWs:
#PMDISK=/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:4:0

#PMMOUNT=/var/lib/pgsql
```

For a nonstandard configuration, modify the persistent disk device names accordingly.

- f. 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. Ignore these particular messages.

8. Reboot the SMW.

```
smw# reboot
```

- (Optional) Set up a virtual network computing (VNC) server for remote access to a desk-side SMW (see [Enable Remote Access to the SMW using VNC](#) on page 147). This is not necessary or recommended for rack-mount SMWs.

**IMPORTANT:** Do not execute any SMW commands until after rebooting the SMW and completing 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, resolve that issue and then 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

#### Procedure

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

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

- RESTRICTION:** This step is 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 and go to step 3 on page 49.

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, be prepared to enter 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. During 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 the user to disable any blades that should not be powered on. After disabling these blades, select 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.

4. **RESTRICTION:** This step is 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 and go to step 5 on page 50.

Complete the bootstrap process on a Cray XC Series system (Aries™ network) SMW:

- a. Power down the system.

```
smw# xtcli power down s0
```

- b. 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
```

- c. Power up the system.

```
smw# xtcli power up s0
```

5. **RESTRICTION:** This step is for a Cray XE or a Cray XK system (Gemini network) SMW only. For Cray XC Systems, the bootstrap process is now complete. Continue with Cray system hardware discovery.

Complete the bootstrap process on a Cray XE or a Cray XK system (Gemini network) SMW.

- a. Exit root.

```
smw# exit
```

- b. 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]$" -f /tmp/xtrsh.Sk5yg9Uz -l root -s "if [-x /etc/
loadnor]; \
then /etc/loadnor; fi ; if md5sum /dev/mtd/4 | grep
fec74d6d797dd8b36e68282ad43d0773; \
then echo OLDBIOS; fi"
There is one L1 not up-to-date.
spawn fm -qRM -t 10 s0
It took 2 seconds for 'query' to complete to all L0s.
xtrsh -m s -f /tmp/xtrsh.6ef31uGH -l root -s "if [-x /etc/loadnor]; \
then /etc/loadnor; fi ; if md5sum /dev/mtd/4 | grep
fec74d6d797dd8b36e68282ad43d0773; \
then echo OLDBIOS; fi"
There are 24 L0s not up-to-date.
#####
Attempt #1
#####
spawn fm -RM -t 11 c0-0
It took 157 seconds (2 minutes, 37 seconds) for 'flash' to complete to all
L1s.
There is one L1 not up-to-date.
xtrsh -m "-[0-7]$" -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 10 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. Continue to [Discover the Cray System Hardware](#) on page 52.
- If the L0s do not get flashed with the updated BIOS firmware, `xtflash` halts and displays 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.



**CAUTION:** Cray recommends reflashing the Cray L0s in groups of a few cabinets at a time rather than reflashing 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 manually. Flashing only a few cabinets at a time limits the risk associated with a catastrophic failure.

**IMPORTANT:** If the BIOS is flashed on an L1/L0 that already has a compatible BIOS (like an SIO), then the L1/L0 must be power-cycled after being flashed.

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
```

With the `fm` command, specific chassis or cabinets can be specified. 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
```

- c. Execute the `xtflash` command again after the BIOS chips have been reflashed, to continue 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

### About this task

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. To display the configuration, use the `xtcli` command after running `xtdiscover`. 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 confirms whether to abort or continue. Provide the X and Y cabinet sizes as well as topology from [Site-dependent Configuration Values](#) in [Configuration Values](#) on page 9 for `xtdiscover` to accurately discover the hardware. The `xtdiscover` command may prompt the user to execute the `xtbounce` command in a separate window and then continue.

The following steps include an example of the output of the `xtdiscover` command.

### Procedure

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

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

2. Run the `xtdiscover` command.

```
smw# xtdiscover
***** xtdiscover started *****
Command line: xtdiscover
USER: crayadm pts/1 Sep 10 07:03 (mycomputer.company.com)

Checking system configuration...

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)[a]: a

Setting system type to Aries
Discovering Aries-based system...

Enter maximum X cabinet size (columns) [1-64, last: 1], q=quit: 1
Enter maximum Y cabinet size (rows) [1-32, last: 1], q=quit: 1
Adding hosts and routes for 1 cabinet...done.

A copy of the current partition configuration has been saved in:
 /home/crayadm/hss_db_backup/partitions.09-10-2015.0728

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

A backup copy of the HSS database has been saved in:
 /home/crayadm/hss_db_backup/db_backup.09-10-2015.0728.sql

A backup copy of the /etc/hosts file has been saved in:
 /home/crayadm/hss_db_backup/hosts.09-10-2015.0728

Checking current configuration...
Component data received from fstream.

Found 0 existing node power management descriptors
done.

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

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

Checking for duplicate MAC addresses...
Done.

Enter your system's network topology class [last: 2]:
Setting topology class to 2

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

Enter the boot node name [c0-0c0s0n1]:
Enter the SDB node name [c0-0c0s0n2]:
Enter the absolute pathname to the default boot image [/raw0]:

Verifying phase 1 (cabinet) configuration...verification complete.

Storing base cabinet data...done.
Resuming State Manager for power-up and bounce...
Resume successful.

Discovery Phase 1 complete.

xtdiscover is about to power on 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]:

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

Loading base component data for discovery phase 2... No updates: skip
done.

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

Discovering component phase 2 (blade) state:
[46 out of 46]
Finished discovering component phase 2 (blade) state.

Discovering component phase 2 (blade) attributes:
[46 out of 46]
Finished discovering component phase 2 (blade) attributes.

Verifying phase 2 (blade) configuration...
INFO: 46 blades(s) were modified (component remains, but with changes).

INFO: Modified the following hardware:
 46 slots

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

verification complete.

Checking for duplicate MAC addresses...
Done.

Summary of blades discovered:
Total: 48 Service: 7 Empty: 2 Disabled: 0

Discovery Phase 2 complete.

Blades needing to be bounced: 46.

Storing attribute data...done.

Resuming State Manager for bounce...
Resume successful.

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

Note that this command will also check for certain types
of BIOS errors in the event of a bounce failure.
If this occurs, please follow any additional directions.

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

### 3. Bounce the system (as `crayadm`) in a separate window.

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

4. After the `xtbounce` command from the previous step has finished, return to the window in which `xtdiscover` was being run to respond to the remaining prompts, as shown below. Type "c" to continue. And then later type "y" when asked to commit the `xtdiscover` results to the database.

```
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:
[46 out of 46]
Finished discovering component phase 3 attributes.
```

```
Discovery Phase 3 complete.
```

```
Verifying final configuration...
```

```
INFO: 112 total component(s) were added (new; not in DB previously, in any
state).
```

```
INFO: 46 total component(s) were modified (component remains, but with changes).
```

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

```
14 pdcs (14 enabled, 0 empty or disabled)
20 accels (20 enabled, 0 empty or disabled)
56 qpdcs (56 enabled, 0 empty or disabled)
10 gpdcs (10 enabled, 0 empty or disabled)
12 hpdcs (12 enabled, 0 empty or disabled)
```

```
INFO: Modified the following hardware:
46 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: xtdiagd xtremoted erfbsd 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 xtdiagd
Flushing and installing cabinet routes...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 *****
```

5. **RESTRICTION:** For a Cray XC30 system (Aries network) SMW only. Skip this step (and its substeps) if installing on a Cray XE or XK system.

Update the firmware for a Cray XC30 system:

- a. Exit from root.

```
smw# exit
```

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

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

If the system was not bounced previously, the following message may be displayed:

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

If so, enter **y**.

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

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



**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.

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

- d. Run `xtzap` with one or more of the options described below to update if the output of `xtzap` from the previous step includes a "Revision Mismatches" section, indicating some firmware is out of date and needs to be reflashed.

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 there is uncertainty about what needs to be updated, use the `-a` option:

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

- e. 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
```

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

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

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

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

6. As user `root`, place the Cray SMW 7.2UP04 Software DVD in the drive and mount it again.

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

7. Execute the `SMWconfig` command to configure MySQL.

When prompted for the old root MySQL database password, press the `Enter` key (equivalent to entering an empty string, which is the default password), and create a new root MySQL database password.

**NOTE:** To become familiar with the requirements for setting and using the root MySQL password, in particular with regard to using the special characters `* ? [ < > & ; ! | $`, review the Oracle [MySQL documentation](#).

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.
```

8. **NOTE:** This step is optional. Perform this step to configure `postfix` on this SMW as a mail transfer agent.

Change the following setting in the `/etc/sysconfig/mail` file on the SMW to prevent the `master.cf` and `main.cf` `postfix` configuration files from being recreated during software updates or fixes.

```
MAIL_CREATE_CONFIG="no"
```

9. Unmount the Cray SMW 7.2UP04 Software DVD.

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

## Confirm the SMW is Communicating with System Hardware

### About this task

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.

### Procedure

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.

**NOTE:** In this example, the `Contents of timestamp...` line appears only in output from Cray XE systems, while output from Cray XC systems contains only the `How long have they been up...` section.

```
crayadm@smw> xtmcinfo -t -u s0
Contents of timestamp...
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
```

## Change the Default System Management Workstation (SMW) Passwords

### About this task

The SMW contains its own `/etc/passwd` file that is separate from the `password` file for the rest of the system. After logging on to the SMW for the first time, it is recommended to change the default passwords, as described in the following instructions:

### Procedure

1. Log on to the SMW as `root`.
2. Execute the following commands:

```
smw# passwd root
smw# passwd crayadm
smw# passwd cray-vnc
smw# passwd mysql
```

For rack-mount SMWs it is also necessary to change the default iDRAC password. For more information, see [Change the Default iDRAC Password](#) on page 60.

## Change the Default iDRAC Password

### About this task

### Procedure

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 is 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.

## Set Up the SUSE Firewall and IP Tables

### About this task

The SMW software includes a firewall. The following steps enable and configure the firewall.

**TIP:** It is not necessary to shut down the system before performing this task.

### Procedure

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. Using the `SuSEfirewall12` program and the following steps, 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 10.3.1.1,udp,514 10.3.1.254,tcp,
6811:6815"
```

**NOTE:** 10.3.1.254 is the boot node's IP address for `eth0` on the network between the boot node and the SMW.

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. Verify the 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 using VNC](#) on page 147.

## Configure the Boot RAID

**NOTE:** Cray ships systems with much of this software installed and configured. Performing all of the steps in the following boot RAID procedures may not be necessary unless the configuration needs to be changed.

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 RAID, see *Manage Lustre for the Cray Linux Environment (CLE)* (S-0010), which is provided with the CLE release package.

Cray provides support for system boot RAID from two different vendors: Data Direct Networks, Inc. (DDN) and NetApp, Inc. To configure the boot RAID, use the procedures for the type of boot RAID at this site:

|                                        |                                                                                                                                                                                                                                                                    |
|----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>DDN boot RAID</b>                   | <p><a href="#">Configure and Zone the Boot RAID for a DDN Storage System</a> on page 63</p> <p>Includes procedures to configure and zone LUNs (logical unit numbers) using either <code>telnet</code> or the Silicon Storage Appliance Manager GUI.</p>            |
| <b>NetApp, Inc. Engenio™ boot RAID</b> | <p><a href="#">Configure the Boot RAID for a NetApp, Inc. Engenio Storage System</a> on page 71</p> <p>Includes procedures to install the SANtricity Storage Manager utility, which is then used to create and configure volumes and then assign them to LUNs.</p> |

After configuring the boot RAID, use the following procedures, as applicable:

- If a QLogic Fibre Channel Switch is used, follow the procedures in [Zone the QLogic™ FC Switch](#) on page 75.

- To rediscover the LUNs (logical unit numbers) and zones that have been created, use this procedure: [Reboot the SMW and Verify LUNs are Recognized](#) on page 78.
- To partition the LUNs, use this procedure: [Partition the LUNs](#) on page 78.
- For systems that use the CLE feature Direct-attached Lustre (DAL), use this procedure: [Create LUNs and an IMPS File System for DAL](#) on page 78.

Prerequisites for these procedures:

- 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 an FC switch.
- The boot nodes have an FC or SAS connection to the boot RAID or to an FC switch.
- The SDB node has an FC or SAS connection to the boot RAID or to an FC switch. The SDB node may have an Ethernet connection.
- If a dedicated `syslog` node is configured, it has an FC or SAS connection to the boot RAID or to an FC switch. A dedicated `syslog` node does not have an Ethernet connection.

## Configure and Zone the Boot RAID for a DDN Storage System

**NOTE:** The instructions in the procedures listed below supersede the documentation supplied by the RAID manufacturer.

- To configure a DDN storage system using the command-line interface (`telnet`), use the following procedures:
  1. [Change the Boot RAID Password for DDN Devices using telnet](#) on page 64
  2. [Configure the LUNs for DDN Devices using telnet](#) on page 65
  3. [Zone the LUNs for DDN Devices using telnet](#) on page 66
- To configure a DDN storage system using the DDN Silicon Storage Appliance (S2A) Manager GUI:
  1. Install the GUI software using these procedures:
    - a. [Install the Silicon Storage Appliance Manager Software for DDN Devices](#) on page 69
    - b. [Identify the Installed Version of the DDN Silicon Storage Appliance Manager Software](#) on page 70
  2. Change the password. The password should be changed on all DDN boot RAIDs, whether using `telnet`, as described in procedure 2 above, or using the S2A Manager GUI. See DDN documentation on the S2A Manager for instructions on how to change the password using the GUI.
  3. Configure and zone LUNs using this procedure: [Configure and Zone the LUNs for DDN Devices using the S2A Manager GUI](#) on page 70

To enable and configure remote logging of DDN messages so that the controller's internal error log does not overflow, use this procedure : [Configure Logging using Syslog on DDN Storage Devices](#) on page 70. This is a `telnet` procedure. DDN documentation for the S2A Manager may have instructions on how to do this using the GUI.

## Change the Boot RAID Password for DDN Devices using `telnet`

### About this task

As part of configuring the boot RAID using DDN, Cray recommends changing the passwords for `admin` and `user`.

The RAID default administrative login name is `admin` and the default password is `password`. The standard IP address for the RAID controller is `10.1.0.1`. The secondary RAID controller, if used, is `10.1.0.2`.

### Procedure

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. Log on 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 using `telnet`

### About this task

Configure the boot RAID with a minimum of six LUNs to support the various system management file systems. Cray recommends nine LUNs, configured as described in "Configuring the Boot RAID LUNs or Volume Groups" in *CLE Installation and Configuration Guide (S-2444)*.

Three LUNs are required for High Availability SMW. If the CLE feature Direct Attached Lustre (DAL) is configured on SMW HA, another LUN is needed for Image Management and Provisioning System (IMPS) data.

### Procedure

1. On the RAID controller, run 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<br>Size | Block<br>Tiers | Tier list |
|-------|-------|-------|--------|----------|------------------|----------------|-----------|
| ----- |       |       |        |          |                  |                |           |

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

```
CAB01-L1: lun delete=x
```

where `x` is the LUN to be deleted.

3. Run the `lun add` 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.

```
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
```

Repeat this step for each LUN to be configured and formatted.

The recommended boot RAID LUN configuration is shown in the following table from *CLE Installation and Configuration Guide (S-2444)*, which is provided with the CLE release package.

Table 8. Recommended Boot RAID LUN Values

| LUN | Label     | Size (1-50 Cabinets) | Size (50+ Cabinets) | Segment Size |
|-----|-----------|----------------------|---------------------|--------------|
| 0   | bootroot0 | 40GB                 | 70GB                | 256KB        |
| 1   | shroot0   | 280GB                | 370GB               | 256KB        |
| 2   | sdb0      | 60GB                 | 80GB                | 256KB        |
| 3   | bootroot1 | 40GB                 | 70GB                | 256KB        |
| 4   | shroot1   | 280GB                | 370GB               | 256KB        |
| 5   | sdb1      | 60GB                 | 80GB                | 256KB        |
| 6   | bootroot2 | 40GB                 | 70GB                | 256KB        |
| 7   | shroot2   | 280GB                | 370GB               | 256KB        |
| 8   | sdb2      | 60GB                 | 80GB                | 256KB        |

**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 using `telnet`

### About this task

After the LUNs are configured and formatted, grant host access to the LUNs 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.

## Procedure

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:
```

3. 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 finished 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! ***
```

| Port | World Wide Name  | LUN Zoning<br>(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! ***
```

| Port | World Wide Name  | LUN Zoning<br>(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! ***
```

```

 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 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 performing maintenance on disk subsystems. Running LUN verification and swap or moving a back-end channel cable could 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
 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 006,006 007,007 008,008
4 24000001FF020320

```

6. When 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 finishing creating, formatting, and zoning the LUNs on the boot RAID, partition them. Although partitioning occurs on the SMW, use the boot LUN partition table included in *CLE Installation and Configuration Guide* (S-2444), which is provided with the CLE release package, because recommendations may change with each CLE release.

## Install the Silicon Storage Appliance Manager Software for DDN Devices

### About this task

This procedure describes how to install the DDN Silicon Storage Appliance (S2A) Manager on the SMW. The S2A Manager has a graphical user interface, which can be used as an alternative to `telnet` to configure and zone the LUNs for DDN storage devices.

### Procedure

1. Log on to the SMW and enable X windows port forwarding.
  - If already logged on to the SMW, `su` to `root` and enter the following command:

```
smw# ssh -X
```

- If not logged on to the SMW, enter the following command:

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

2. Copy (as `root` user) the Silicon Storage Appliance Manager installation file (`directGUI_version_Linux_x64_NoVM.bin`) to a temporary location on the SMW, such as `/tmp`.  
The Silicon Storage Appliance Manager installation file is on a CD that is provided with the system.
3. Move to the directory that the installation file was copied to.

```
smw# cd /tmp
```

4. Initiate the installation process:

```
smw# ./directGUI_version_Linux_NoVM.bin
```

The **Introduction** window displays.

5. Select **Next**.  
The **License Agreement** window displays.
6. Accept the License Agreement and select **Next**.  
The **Choose Uninstall Folder** window displays, showing the installed software.
7. Select the uninstall folder (file path) to uninstall and select **Next**.
  - a. Select the **complete** icon for a complete uninstall.
  - b. Select **Done**.The **Choose Java Virtual Machine** window opens.
8. Choose a Java Virtual Machine and select **Next**.

The **Choose Install Folder** window opens.

9. Choose the install folder and select **Next**.

The **Choose Link Folder** window opens.

10. Choose the link folder and select **Next**.

The **Pre-Installation Summary** window opens.

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

When the installation completes, the **Install Complete** window opens, indicating the installation is complete.

12. Select **Done**.

## Identify the Installed Version of the DDN Silicon Storage Appliance Manager Software

### Prerequisites

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

### About this task

### Procedure

1. Start the GUI using one of the following:

- `smw# /usr/local/ddn/Silicon_Storage_Appliance_Manager`
- `smw# /opt/Silicon_Storage_Appliance_Manager/Silicon_Storage_Appliance_Manager`

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

## Configure and Zone the LUNs for DDN Devices using the S2A Manager GUI

Use the DDN Silicon Storage Appliance (S2A) Manager to configure and zone LUNs. See the DDN S2A Manager GUI documentation for instructions.

## Configure Logging using Syslog on DDN Storage Devices

### About this task

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 reconfiguring the S2A network logging to enable the SMW to store and retrieve failure information.

When a RAID controller is configured 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 use an IP address beginning with `10.1.0`. If the RAID controllers begin with a different IP address, set the `llm_raid_ip`

variable in the `SMWinstall.conf` file, remount the SMW installation media, and rerun the SMW installer with the `forceupdate` option.

```
smw# vi /home/crayadm/SMWinstall.conf
smw# mount -o loop,ro smw-image-7.2.0-1.0702.37275.648-1.iso /media/cdrom
smw# /media/cdrom/SMWinstall --forceupdate
```

## Procedure

### 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 was set in [Change the Boot RAID Password for DDN Devices using `telnet`](#) on page 64.

```
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 IP address `10.1.1.1` for the SMW `eth1` interface 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 for a NetApp, Inc. Engenio Storage System

To configure the boot RAID for a NetApp, Inc. Engenio Storage System, it is necessary to first install the SANtricity Storage Manager Utility. Then, the rest of the procedures use that utility to create and configure volumes, assign them to LUNs, and configure remote logging.

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

1. Install the SANtricity Storage Manager Utility using this procedure: [Install SANtricity Storage Manager Software for NetApp, Inc. Engenio Devices](#) on page 72. SANtricity is provided as a separate package and is installed from a CD. It may already be installed on the SMW.
2. Use the SANtricity Storage Manager utility from NetApp, Inc. to perform the remaining procedures. These procedures assume familiarity with using the SANtricity interface.
  - a. To create the volume group, use this procedure: [Create the Boot RAID Volume Group for NetApp, Inc. Engenio devices](#) on page 73.
  - b. To create the LUNs within the volume group, use this procedure: [Create and Configure Volumes for NetApp, Inc. Engenio Devices](#) on page 74.
  - c. To configure remote logging of the boot RAID messages, use this procedure: [Configure Remote Logging of NetApp, Inc. Engenio Storage System Boot RAID Messages](#) on page 75.

## 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. If the SANtricity software is installed, then the `SMclient` executable will be found in `/opt/SMgr/client/SMclient`. If this Cray system does not have the software installed on the SMW, then install it using the procedure in [Install the SANtricity Software](#) on page 72.

### Install the SANtricity Software

## Procedure

1. Log on to the SMW as `root`.

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

2. Install SANtricity Storage Manager from the CD or from a directory.
  - To install 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
```

- To install 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 installing from the CD:

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

If 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, 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
```

## Create the Boot RAID Volume Group for NetApp, Inc. Engenio devices

### About this task

Create the 3+1 Volume Group and 1 Global Hot Spare across the first five disks for a 4.5 TB Volume Group (the amount of storage for this installation may be different). The **Array Management** window should still be displayed after performing the procedure.

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

### Procedure

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

**NOTE:** The following steps assume the **Manual** option was selected.
3. Double-click the name for the storage array that to be configured. 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.

After creating the first **Volume Group**, create the first volume when prompted.

## Create and Configure Volumes for NetApp, Inc. Engenio Devices

### About this task

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

Three LUNs are required for High Availability SMW. If the CLE feature Direct Attached Lustre (DAL) is configured on SMW HA, another LUN is needed for Image Management and Provisioning System (IMPS) data.

Use the boot LUN configuration table included in *CLE Installation and Configuration Guide (S-2444)*, which is provided with the 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.

### Procedure

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 to be created. If this is the last volume, select **No** and skip to step 14 on page 75
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 *CLE Installation and Configuration Guide (S-2444)*, which is provided with the 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 the NetApp, Inc. Engenio Storage System documentation for additional information.

### Zone the QLogic™ FC Switch

For a QLogic Fibre Channel Switch, follow [Configure Zoning for a QLogic SANbox Switch Using QuickTools Utility](#) on page 76 to zone the LUNs on the 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. It requires 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

### Procedure

1. Start a web browser.
2. Enter the IP address of the switch. If the 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 the browser. Select **Add Fabric**.  
**NOTE:** If a dialog box appears stating 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 creating a backup of the switch configuration ([Create a Backup of the QLogic Switch Configuration](#) on page 77) before closing and exiting the application.

## Create a Backup of the QLogic Switch Configuration

### About this task

Use the QuickTools utility to create a backup of the QLogic switch configuration. To use QuickTools, a Java browser plugin, version 1.4.2 or later is required.

To start a web browser and open the QuickTools utility, complete steps 1 through 4. If the QuickTools utility is already open, skip to step 5.

### Procedure

1. Start a web browser.
2. Enter the IP address of the 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. If the configuration has a single switch, the IP address is 10.1.0.250.
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 a dialog box appears stating 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 the QLogic switch configuration.
  - c. Select the **Save** button.
6. Close and exit the application.

## Reboot the SMW and Verify LUNs are Recognized

### About this task

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

### Procedure

1. Log on as the `root` user.

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

2. Reboot the SMW 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 *CLE Installation and Configuration Guide (S-2444)*, which is provided with the CLE release package.

```
smw# fdisk -l
```

### Partition the LUNs

After creating, formatting, and zoning the LUNs on the boot RAID, partition them. For procedures and LUN partitioning recommendations, see *CLE Installation and Configuration Guide (S-2444)*, which is provided with the CLE release package.

### Set Up Boot RAID Space for Direct-attached Lustre

If the system will use Direct-attached Lustre (DAL), 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

### Procedure

1. Follow the instructions in [Configure the Boot RAID](#) on page 62 and [Configure and Zone the Boot RAID for a DDN Storage System](#) on page 63 to create and configure LUNs on the 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
Command (m for help): 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>
```

## Configure the Simple Event Correlator (SEC)

The System Management Workstation (SMW) 7.2.UP04 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.

To use the Cray SEC, see *Configure SEC Software* (S-2542) for configuration procedures.

## Back Up the Newly Installed and Configured SMW Software

After installing and configuring the new SMW software, create a backup of it and, if desired, set up the bootable backup drive as an alternate boot device.

For an **R630 rack-mount SMW**, use the following two procedures:

1. [R630 SMW: Create an SMW Bootable Backup Drive](#) on page 80
2. [\(Optional\) R630 SMW: Set Up the Bootable Backup Drive as an Alternate Boot Device](#) on page 88

For an **R815 rack-mount SMW**, use the following two procedures:

1. [R815 SMW: Create an SMW Bootable Backup Drive](#) on page 90
2. [\(Optional\) R815 SMW: Set Up the Bootable Backup Drive as an Alternate Boot Device](#) on page 98

For a **desk-side SMW**, use the following two procedures:

1. [Desk-side SMW: Create an SMW Bootable Backup Drive](#) on page 100
2. [\(Optional\) Desk-side SMW: Set Up the Bootable Backup Drive as an Alternate Boot Device](#) on page 106

### R630 SMW: Create an SMW Bootable Backup Drive

#### About this task

This procedure creates a bootable backup drive for a Dell R630 SMW in order to replace the primary drive if the primary drive fails. When these steps are completed, the backup drive on the SMW will be a bootable replacement for the primary drive when the backup drive is plugged in as or cabled as the primary drive.

**IMPORTANT:** The disk device names shown in this procedure are only examples. Substitute the actual disk device names for the actual system. The boot disk is `pci-0000:03:00.0-scsi-0:0:0:0` and is slot 0, and the bootable backup disk is `pci-0000:03:00.0-scsi-0:0:1:0` and is slot 1.

**NOTICE:** To create a clean backup, Cray recommends shutting down the Cray system before beginning this procedure.

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

#### Procedure

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

```
crayadm@smw> su -
Password:
smw#
```

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 this site has completed this step at least once after installing the updated SMW base



operating system), skip to step 3 on page 83; 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 the R630 SMW.

- a. Identify the installed SMW drive model numbers, serial numbers, and associated Linux device (/dev) names.

Execute `smwmapdrives` 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 running that release or a later one, 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/scsi-35000c50079ab34b7
 /dev/disk/by-id/wwn-0x5000c50079ab34b7
 /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0
Physical slot 1:
 /dev/sdb
 /dev/disk/by-id/scsi-35000c50079ab71c4
 /dev/disk/by-id/wwn-0x5000c50079ab71c4
 /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0
Physical slot 2:
 /dev/sdc
 /dev/disk/by-id/scsi-35000c50079ab313b
 /dev/disk/by-id/wwn-0x5000c50079ab313b
 /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:2:0
Physical slot 3:
 /dev/sdd
 /dev/disk/by-id/scsi-35000c50079ab4b4c
 /dev/disk/by-id/wwn-0x5000c50079ab4b4c
 /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:3:0
Physical slot 4:
 /dev/sde
 /dev/disk/by-id/scsi-35000c50079d05e70
 /dev/disk/by-id/wwn-0x5000c50079d05e70
 /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:4:0
Physical slot 5:
 NOT INSTALLED
Physical slot 6:
 NOT INSTALLED
Physical slot 7:
 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; these 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 desk-side 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
```

Cray recommends that `/boot/grub/device.map`, `/etc/fstab` and `/boot/grub/menu.lst` changes use the "by-path" rather than the "by-id" device name because that would allow physically swapping the backup drive into the primary slot when there is a disk failure in the primary disk. If the backup disk is intended as backup only, rather than as a bootable backup, it is acceptable to use either device name.

- 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 r630 SMW
grub(8) device mapping for boot-drive identification
hd? numbers are being mapped to their physical
(hd0) /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0
(hd1) /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0
(hd2) /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:2:0
(hd3) /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:3:0
(hd4) /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:4:0
(hd5) /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:5:0
```

- 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 step 2.a on page 81 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`.

For example, if the `root` and `swap` file systems are currently configured to mount `/dev/sda2`, they should be changed. Using the `by-path` device name from the example in step 2.a on page 90, the `fstab` lines would change from:

```
/dev/sda1 swap swap defaults 0 0
/dev/sda2 / ext3 acl,user_xattr 1 1
```

to:

```
/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0-part1 swap swap
defaults 0 0
/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0-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-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-0000:03:00.0-scsi-0:0:0:0
Disk /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0: 500.1 GB, 500107862016 bytes
255 heads, 63 sectors/track, 60801 cylinders, total 976773168 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x000c3cc8
```

| Blocks                                                | Id | System               | Device | Boot | Start    | End       |
|-------------------------------------------------------|----|----------------------|--------|------|----------|-----------|
| /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0-part1 |    |                      |        |      | 2048     | 67102719  |
| 33550336                                              | 82 | Linux swap / Solaris |        |      |          |           |
| /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0-part2 | *  |                      |        |      | 67102720 | 976773119 |
| 454835200                                             | 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 the changes to the disk. For help, type `m` within `fdisk`.

```
smw# fdisk -u /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0

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
Partition number (1-4, default 1): 1
First sector (2048-976773167, default 2048): (Press the Enter key)
Using default value 2048
Last sector, +sectors or +size{K,M,G} (2048-976773167, default 976773167): 67102719

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, default 2): 2
First sector (67102720-976773167, default 67102720): (Press the Enter key)
Using default value 67102720
Last sector, +sectors or +size{K,M,G} (67102720-976773167, default 976773167): (Press the Enter key)
Using default value 976773167

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 the `pci-0000:03:00.0-scsi-0:0:0:0` sector information.

```
smw# fdisk -lu /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0

Disk /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0: 500.1 GB, 500107862016 bytes
255 heads, 63 sectors/track, 60801 cylinders, total 976773168 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x7c334e96
```

| Blocks                                                | Id | System               | Device | Boot | Start    | End       |
|-------------------------------------------------------|----|----------------------|--------|------|----------|-----------|
| /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0-part1 |    |                      |        |      | 2048     | 67102719  |
| 33550336                                              | 82 | Linux swap / Solaris |        |      |          |           |
| /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0-part2 |    |                      |        |      | 67102720 | 976773167 |
| 454835224                                             | 83 | Linux                |        |      |          |           |

4. Initialize the swap device.

```
smw# mkswap /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0-part1
Setting up swapspace version 1, size = 33550332 KiB
no label, UUID=8391498b-d159-469c-b766-66f00a28ff74
```

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-0000:03:00.0-scsi-0:0:1:0-part2
mke2fs 1.41.9 (22-Aug-2009)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
28434432 inodes, 113708806 blocks
5685440 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=4294967296
3471 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, 78675968,
 102400000

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.
```

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

```
smw# mount /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0-part2 /mnt
```

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

```
smw# df
Filesystem 1K-blocks Used Available Use% Mounted on
/dev/sda2 447696736 9180648 437606420 3% /
udev 66029308 744 66028564 1% /dev
tmpfs 66029308 39540 65989768 1% /dev/shm
/dev/sdae 309637120 1107516 292800964 1% /var/opt/cray/disk/1
```

```

/dev/sdac 206424760 1963664 193975336 2% /home
/dev/sdad 154818540 474696 146479524 1% /var/lib/mysql
/dev/drbd_r0 961405840 247180 912322076 1% /var/lib/pgsql
/dev/sdb2 447696760 202940 424752060 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-0000:03:00.0-scsi-0:0:0:0-part2 | restore rf -
DUMP: WARNING: no file `/etc/dumpdates'
DUMP: Date of this level 0 dump: Wed Sep 16 15:40:41 2015
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 9129804 blocks.
DUMP: Volume 1 started with block 1 at: Wed Sep 16 15:43:08 2015
DUMP: dumping (Pass III) [directories]
DUMP: dumping (Pass IV) [regular files]
restore: ./lost+found: File exists
./tmp/rstdir1442436041: (inode 27254928) not found on tape
./tmp/rstmodel442436041: (inode 27254931) not found on tape
DUMP: 77.64% done at 23626 kB/s, finished in 0:01
DUMP: Volume 1 completed at: Wed Sep 16 15:50:09 2015
DUMP: Volume 1 9132800 blocks (8918.75MB)
DUMP: Volume 1 took 0:07:01
DUMP: Volume 1 transfer rate: 21693 kB/s
DUMP: 9132800 blocks (8918.75MB)
DUMP: finished in 421 seconds, throughput 21693 kBytes/sec
DUMP: Date of this level 0 dump: Wed Sep 16 15:40:41 2015
DUMP: Date this dump completed: Wed Sep 16 15:50:09 2015
DUMP: Average transfer rate: 21693 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.

```
smw# vi /mnt/etc/fstab
```

For example, change

```

/dev/disk/by-id/scsi-35000c50079ab34b7-part1 swap swap
defaults 0 0
/dev/disk/by-id/scsi-35000c50079ab34b7-part2 / ext3
acl,user_xattr 1 1

```

to:

```

/dev/disk/by-id/scsi-35000c50079ab71c4-part1 swap swap
defaults 0 0

```

```
/dev/disk/by-id/scsi-35000c50079ab71c4-part2 / ext3
acl,user_xattr 1 1
```

- 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/scsi-35000c50079ab34b7-part2
```

with:

```
root=/dev/disk/by-id/scsi-35000c50079ab71c4-part2
```

or 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` detects only the first 16 devices. Therefore, if a disk is added 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_1
```

- 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_1
find /THIS_IS_1
(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
quit
```

**IMPORTANT:** For R630 SMWs, grub recreates `device.map` with the short names, not the persistent names. Do not trust the `/dev/sdx` names. Always use `find` when executing grub because it is possible that grub root may not be `hd2` the next time grub is executed.

## 11. Unmount the backup root partition.

```
smw# umount /mnt
```

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

## (Optional) R630 SMW: Set Up the Bootable Backup Drive as an Alternate Boot Device

### About this task

This optional procedure modifies a bootable backup drive for a Dell R630 SMW in order to boot from and run the R630 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. Substitute the actual disk device names for this system. The boot disk is `pci-0000:03:00.0-scsi-0:0:0:0` and is slot 0, and the bootable backup disk is `pci-0000:03:00.0-scsi-0:0:1:0` and is slot 1.

## Procedure

### 1. Mount the backup drive's root partition.

```
smw# mount /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0-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:

- 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.
- 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 `pci-0000:03:00.0-scsi-0:0:0:0` (slot 0) as the primary drive and `pci-0000:03:00.0-scsi-0:0:1:0` (slot 1) as the backup drive.

```
smw# cp -p /boot/grub/menu.lst /boot/grub/menu.lst.20150916
smw# vi /boot/grub/menu.lst
smw# cat /boot/grub/menu.lst
```



```
Modified by YaST2. Last modification on Thu Aug 13 19:38:47 CDT 2015
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 - 3.0.101-0.46
 root (hd0,1)
 kernel /boot/vmlinuz-3.0.101-0.46-default root=/dev/disk/by-path/
pci-0000:03:00.0-scsi-0:0:0:0-part2 pci=bfsort resume=/dev/disk/by-path/
pci-0000:03:00.0-scsi-0:0:0:0-part1 splash=silent crashkernel=256M-:128M@16M
showopts biosdevname=X vga=0x31a
 initrd /boot/initrd-3.0.101-0.46-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 - 3.0.101-0.46
 root (hd1,1)
 kernel /boot/vmlinuz-3.0.101-0.46-default root=/dev/disk/by-path/
pci-0000:03:00.0-scsi-0:0:1:0-part2 pci=bfsort resume=/dev/disk/by-path/
pci-0000:03:00.0-scsi-0:0:1:0-part1 splash=silent crashkernel=256M-:128M@16M
showopts biosdevname=X vga=0x31a
 initrd /boot/initrd-3.0.101-0.46-default

###Don't change this comment - YaST2 identifier: Original name: failsafe###
title Failsafe -- SUSE Linux Enterprise Server 11 SP3 - 3.0.101-0.46
 root (hd0,1)
 kernel /boot/vmlinuz-3.0.101-0.46-default root=/dev/disk/by-path/
pci-0000:03:00.0-scsi-0:0:0:0-part2 showopts ide=nodma apm=off noresume edd=off
powersaved=off nohz=off highres=off processor.max_cstate=1 nomodeset
x11failsafe biosdevname=X vga=0x31a
 initrd /boot/initrd-3.0.101-0.46-default
```

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

Examine the backup drive's `fstab` file. Edit the `/mnt/etc/fstab` file, changing

`pci-0000:03:00.0-scsi-0:0:0:0` to `pci-0000:03:00.0-scsi-0:0:1:0` device names to reference the backup drive. In the following example, the backup drive is `pci-0000:03:00.0-scsi-0:0:1:0-...`

```
smw# cp /mnt/etc/fstab /mnt/etc/fstab.20150916
smw# cat /mnt/etc/fstab
/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0-part1swap swap
defaults 0 0
/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0-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
/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:3:0-part1/var/opt/cray/disk/1 ext3
defaults 1 0
none /var/lib/dhcp/db ramfs defaults 0 0

smw# vi /mnt/etc/fstab
smw# cat /mnt/etc/fstab
/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0-part1swap swap
```

```

defaults 0 0
/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0-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
/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:3:0-part1/var/opt/cray/disk/1 ext3
defaults 1 0
none /var/lib/dhcp/db ramfs defaults 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.

## R815 SMW: Create an SMW Bootable Backup Drive

### About this task

This procedure creates a bootable backup drive for a Dell R815 SMW in order to replace the primary drive if the primary drive fails. When these steps are completed, the backup drive on the SMW will be a bootable replacement for the primary drive when the backup drive is plugged in as or cabled as the primary drive.

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

**NOTICE:** To create a clean backup, Cray recommends shutting down the Cray system before beginning this procedure.

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

## Procedure

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

```

crayadm@smw> su -
Password:
smw#

```

### 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 this site has completed this step at least once after installing the updated SMW base operating system), skip to step 3 on page 93; 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 the R815 SMW.

- a. Identify the installed SMW drive model numbers, serial numbers, and associated Linux device (`/dev`) names.

Execute `smwmapdrives` 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 running that release or a later one, 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; these 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 desk-side 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
```

Cray recommends that `/boot/grub/device.map`, `/etc/fstab` and `/boot/grub/menu.lst` changes use the "by-path" rather than the "by-id" device name because that would allow physically swapping the backup drive into the primary slot when there is a disk failure in the primary disk. If the backup disk is intended as backup only, rather than as a bootable backup, it is acceptable to use either device name.

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

```
Dell Rackmount r815 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 step 2.a on page 90 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`.

For example, if the `root` and `swap` file systems are currently configured to mount `/dev/sda2`, they should be changed. Using the `by-path` device name from the example in step 2.a on page 90, the `fstab` lines would change from:

|                        |                   |                   |                             |                  |
|------------------------|-------------------|-------------------|-----------------------------|------------------|
| <code>/dev/sda1</code> | <code>swap</code> | <code>swap</code> | <code>defaults</code>       | <code>0 0</code> |
| <code>/dev/sda2</code> | <code>/</code>    | <code>ext3</code> | <code>acl,user_xattr</code> | <code>1 1</code> |

to:

```
/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
```

- 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-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 the 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 the `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-part2

```



or 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` detects only the first 16 devices. Therefore, if a disk is added 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 R815 SMWs, `grub` recreates `device.map` with the short names, not the persistent names. Do not trust the `/dev/sdx` names. Always use `find` when executing `grub` because it is possible that `grub root` may not be `hd2` the next time `grub` is executed.

11. Unmount the backup root partition.

```
smw# umount /mnt
```

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

## (Optional) R815 SMW: Set Up the Bootable Backup Drive as an Alternate Boot Device

### About this task

This optional procedure modifies a bootable backup drive for a Dell R815 SMW in order to boot from and run the R815 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. Substitute the actual disk device names for this system. The boot disk is `phy7` and is slot 0, and the bootable backup disk is `phy6` and is slot 1.

### Procedure

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.

Examine the backup drive's `fstab` file. 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`-....

```

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
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

### About this task

This procedure creates a System Management Workstation (SMW) bootable backup drive for a desk-side SMW. Its purpose is to replace the primary drive if the primary drive fails. When this procedure is completed, the backup drive on the SMW will be a bootable replacement for the primary drive when the backup drive is plugged in as or cabled as the primary drive.

**IMPORTANT:** The disk device names shown in this procedure are only examples; substitute the actual disk device names when executing this procedure. 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`.

**NOTICE:** To create a clean backup, Cray recommends shutting down the Cray system before beginning this procedure.

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

### Procedure

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 the 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. Be aware that errors in the `/boot/grub/menu.lst` will affect the 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 - 3.0.101-0.46
root (hd0,1)
kernel /boot/vmlinuz-3.0.101-0.46-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-3.0.101-0.46-default
###Don't change this comment - YaST2 identifier: Original name: failsafe###
title Failsafe -- SUSE Linux Enterprise Server 11 - 3.0.101-0.46
root (hd0,1)
kernel /boot/vmlinuz-3.0.101-0.46-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-3.0.101-0.46-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 the 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 - 3.0.101-0.46
kernel /boot/vmlinuz-3.0.101-0.46-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 - 3.0.101-0.46
kernel /boot/vmlinuz-3.0.101-0.46-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 - 3.0.101-0.46
kernel /boot/vmlinuz-3.0.101-0.46-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 - 3.0.101-0.46
kernel /boot/vmlinuz-3.0.101-0.46-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 - 3.0.101-0.46
```

```

root (hd0,1)
kernel /boot/vmlinuz-3.0.101-0.46-default root=/dev/sda2 \
resume=/dev/sda1 splash=silent crashkernel=256M-:128M@16M showopts
initrd /boot/initrd-3.0.101-0.46-default
###Don't change this comment - YaST2 identifier: Original name: failsafe###
title Failsafe -- SUSE Linux Enterprise Server 11 - 3.0.101-0.46
root (hd0,1)
kernel /boot/vmlinuz-3.0.101-0.46-default \
root=/dev/sda2 showopts ide=nodma apm=off noresume edd=off
powersaved=off nohz=off highres=off initrd /boot/initrd-3.0.101-0.46-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 a disk is added 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.

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

### About this task

This optional 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.

The disk device names shown in this procedure are provided as examples only. Substitute the correct disk devices for this 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`.

### Procedure

#### 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. Find the `grub` name of the backup drive 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 - 3.0.101-0.46
 root (hd0,1)
 kernel /boot/vmlinuz-3.0.101-0.46-default root=/dev/sda2 \
 resume=/dev/sda1 splash=silent crashkernel=256M-:128M@16M showopts
vga=0x31a \
 initrd /boot/initrd-3.0.101-0.46-default

###Don't change this comment - YaST2 identifier: Original name: failsafe###
title Failsafe -- SUSE Linux Enterprise Server 11 - 3.0.101-0.46
 root (hd0,1)
 kernel /boot/vmlinuz-3.0.101-0.46-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-3.0.101-0.46-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 - 3.0.101-0.46
 root (hd0,1)
 kernel /boot/vmlinuz-3.0.101-0.46-default root=/dev/sdc2 \
 resume=/dev/sdc1 splash=silent crashkernel=256M-:128M@16M showopts
vga=0x31a \
 initrd /boot/initrd-3.0.101-0.46-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.

## Update or Upgrade the 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 minor version (numbers following *UP*) will change.

Follow the procedures in this chapter to install an SMW 7.2UP04 package. The procedures provided in this chapter do not change the base operating system version running on the SMW.



**CAUTION:** The SMW 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 in order to perform the following update/upgrade procedures.

### Prepare to Upgrade or Update SMW Software

**IMPORTANT:** For a system configured for SMW high availability (HA) with the SMW failover feature, prepare both SMWs for an upgrade.

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

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

- Read the *SMW README* and *SMW Errata* provided in the SMW update directory for any changes to upgrade or update procedures.
- Read the Field Notices (FN) related to kernel security fixes, and apply any needed changes before continuing with the installation.
- If local changes have been made 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 changes that should be merged into the locally modified file.
- If local changes have been made 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: To retain any power management profiles that were created, 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 the system. This file, and a file containing the default values, `hss.ini.dist` are located in `/opt/tftpboot/ccrd`.
  - If `hss.ini` does not have local changes, delete this file before beginning the update installation, and an active `hss.ini` file will be created as part of the installation.
  - If `hss.ini` has local changes, save a local copy of `hss.ini` to another location. After installation is complete, copy `hss.ini.dist` to `hss.ini` and re-create the local changes in the new `hss.ini` file.
- If using the Cray simple event correlator (SEC) and the `/opt/cray/default/SEC_VARIABLES` file has local changes, make a backup copy of this file before beginning the upgrade or update. For more information, see *Configure Cray SEC Software (S-2542)*.
- If `/etc/syslog-ng/syslog-ng.conf` or `/etc/rsyslog.conf` has local changes, the local changes are saved during the upgrade procedure. Also, if `/opt/cray/hss/default/etc/xtnlrd.ini` has local changes and the new release includes an updated `xtnlrd.ini` file, the local 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 the local `xtnlrd.ini` file.
- Update the `properties.local` file. (See [Update the properties.local File](#) on page 110)
- Back up the current software. (See [Back Up the Current Software](#) on page 110).

## Update the properties.local File

If a software update includes changes to the default properties file, the existing `properties.local` file is renamed `/opt/cray/default/pm/properties.local.YYYY-MM-DD.HH:MM:SS`. Any site-specific changes, such as power management profiles, that existed prior to the update will need to be merged into the newly created `/opt/cray/default/pm/properties.local` file.

## Back Up the Current Software

Before installing the new SMW software, back up the current SMW software installation.

- For a Dell R630 (rack-mount) SMW, use [R630 SMW: Create an SMW Bootable Backup Drive](#) on page 80.
- For a Dell R815 (rack-mount) SMW, use [R815 SMW: Create an SMW Bootable Backup Drive](#) on page 90.
- For a desk-side SMW, use [Desk-side SMW: Create an SMW Bootable Backup Drive](#) on page 100.

## R630 SMW: Create an SMW Bootable Backup Drive

### About this task

This procedure creates a bootable backup drive for a Dell R630 SMW in order to replace the primary drive if the primary drive fails. When these steps are completed, the backup drive on the SMW will be a bootable replacement for the primary drive when the backup drive is plugged in as or cabled as the primary drive.

**IMPORTANT:** The disk device names shown in this procedure are only examples. Substitute the actual disk device names for the actual system. The boot disk is `pci-0000:03:00.0-scsi-0:0:0:0` and is slot 0, and the bootable backup disk is `pci-0000:03:00.0-scsi-0:0:1:0` and is slot 1.

**NOTICE:** To create a clean backup, Cray recommends shutting down the Cray system before beginning this procedure.

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

## Procedure

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

```
crayadm@smw> su -
Password:
smw#
```

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 this site has completed this step at least once after installing the updated SMW base operating system), skip to step 3 on page 114; 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 the R630 SMW.

- a. Identify the installed SMW drive model numbers, serial numbers, and associated Linux device (`/dev`) names.

Execute `smwmapdrives` 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 running that release or a later one, 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/scsi-35000c50079ab34b7
/dev/disk/by-id/wwn-0x5000c50079ab34b7
/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0
Physical slot 1:
/dev/sdb
/dev/disk/by-id/scsi-35000c50079ab71c4
/dev/disk/by-id/wwn-0x5000c50079ab71c4
/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0
Physical slot 2:
/dev/sdc
/dev/disk/by-id/scsi-35000c50079ab313b
/dev/disk/by-id/wwn-0x5000c50079ab313b
/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:2:0
Physical slot 3:
/dev/sdd
/dev/disk/by-id/scsi-35000c50079ab4b4c
/dev/disk/by-id/wwn-0x5000c50079ab4b4c
/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:3:0
Physical slot 4:
/dev/sde
/dev/disk/by-id/scsi-35000c50079d05e70
/dev/disk/by-id/wwn-0x5000c50079d05e70
/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:4:0
Physical slot 5:
NOT INSTALLED
Physical slot 6:
NOT INSTALLED
```

```
Physical slot 7:
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; these 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 desk-side 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
```

Cray recommends that `/boot/grub/device.map`, `/etc/fstab` and `/boot/grub/menu.lst` changes use the "by-path" rather than the "by-id" device name because that would allow physically swapping the backup drive into the primary slot when there is a disk failure in the primary disk. If the backup disk is intended as backup only, rather than as a bootable backup, it is acceptable to use either device name.

- 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 r630 SMW
grub(8) device mapping for boot-drive identification
hd? numbers are being mapped to their physical
(hd0) /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0
(hd1) /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0
(hd2) /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:2:0
(hd3) /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:3:0
(hd4) /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:4:0
(hd5) /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:5:0
```

- 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 step 2.a on page 111 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`.

For example, if the `root` and `swap` file systems are currently configured to mount `/dev/sda2`, they should be changed. Using the `by-path` device name from the example in step 2.a on page 90, the `fstab` lines would change from:

|                        |                   |                   |                             |                  |
|------------------------|-------------------|-------------------|-----------------------------|------------------|
| <code>/dev/sda1</code> | <code>swap</code> | <code>swap</code> | <code>defaults</code>       | <code>0 0</code> |
| <code>/dev/sda2</code> | <code>/</code>    | <code>ext3</code> | <code>acl,user_xattr</code> | <code>1 1</code> |

to:

|                                                                    |                   |                   |
|--------------------------------------------------------------------|-------------------|-------------------|
| <code>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0-part1</code> | <code>swap</code> | <code>swap</code> |
| <code>defaults</code>                                              | <code>0 0</code>  |                   |
| <code>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0-part2</code> | <code>/</code>    | <code>ext3</code> |
| <code>acl,user_xattr</code>                                        | <code>1 1</code>  |                   |

- 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-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-0000:03:00.0-scsi-0:0:0:0
Disk /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0: 500.1 GB, 500107862016 bytes
255 heads, 63 sectors/track, 60801 cylinders, total 976773168 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x000c3cc8
```

| Blocks                                                | Id | System               | Device | Boot | Start    | End       |
|-------------------------------------------------------|----|----------------------|--------|------|----------|-----------|
| /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0-part1 |    |                      |        |      | 2048     | 67102719  |
| 33550336                                              | 82 | Linux swap / Solaris |        |      |          |           |
| /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0-part2 | *  |                      |        |      | 67102720 | 976773119 |
| 454835200                                             | 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 the changes to the disk. For help, type `m` within `fdisk`.

```
smw# fdisk -u /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0

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
Partition number (1-4, default 1): 1
First sector (2048-976773167, default 2048): (Press the Enter key)
Using default value 2048
Last sector, +sectors or +size{K,M,G} (2048-976773167, default 976773167): 67102719

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, default 2): 2
First sector (67102720-976773167, default 67102720): (Press the Enter key)
Using default value 67102720
Last sector, +sectors or +size{K,M,G} (67102720-976773167, default 976773167): (Press
```

the Enter key)

Using default value 976773167

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 the `pci-0000:03:00.0-scsi-0:0:0:0` sector information.

```
smw# fdisk -lu /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0
```

```
Disk /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0: 500.1 GB, 500107862016 bytes
255 heads, 63 sectors/track, 60801 cylinders, total 976773168 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x7c334e96
```

| Blocks                                                | Id | System               | Device | Boot | Start    | End       |
|-------------------------------------------------------|----|----------------------|--------|------|----------|-----------|
| /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0-part1 |    |                      |        |      | 2048     | 67102719  |
| 33550336                                              | 82 | Linux swap / Solaris |        |      |          |           |
| /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0-part2 |    |                      |        |      | 67102720 | 976773167 |
| 454835224                                             | 83 | Linux                |        |      |          |           |

4. Initialize the swap device.

```
smw# mkswap /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0-part1
```

Setting up swapspace version 1, size = 33550332 KiB

no label, UUID=8391498b-d159-469c-b766-66f00a28ff74

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-0000:03:00.0-scsi-0:0:1:0-part2
```

mke2fs 1.41.9 (22-Aug-2009)

Filesystem label=

OS type: Linux

Block size=4096 (log=2)

Fragment size=4096 (log=2)

28434432 inodes, 113708806 blocks

5685440 blocks (5.00%) reserved for the super user

First data block=0

Maximum filesystem blocks=4294967296

3471 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, 78675968,
102400000
```

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.

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

```
smw# mount /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0-part2 /mnt
```

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

```
smw# df
Filesystem 1K-blocks Used Available Use% Mounted on
/dev/sda2 447696736 9180648 437606420 3% /
udev 66029308 744 66028564 1% /dev
tmpfs 66029308 39540 65989768 1% /dev/shm
/dev/sdae 309637120 1107516 292800964 1% /var/opt/cray/disk/1
/dev/sdac 206424760 1963664 193975336 2% /home
/dev/sdad 154818540 474696 146479524 1% /var/lib/mysql
/dev/drbd_r0 961405840 247180 912322076 1% /var/lib/pgsql
/dev/sdb2 447696760 202940 424752060 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-0000:03:00.0-scsi-0:0:0:0-part2 | restore rf -
DUMP: WARNING: no file `/etc/dumpdates'
DUMP: Date of this level 0 dump: Wed Sep 16 15:40:41 2015
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 9129804 blocks.
DUMP: Volume 1 started with block 1 at: Wed Sep 16 15:43:08 2015
DUMP: dumping (Pass III) [directories]
DUMP: dumping (Pass IV) [regular files]
restore: ./lost+found: File exists
./tmp/rstidir1442436041: (inode 27254928) not found on tape
./tmp/rstmodel1442436041: (inode 27254931) not found on tape
DUMP: 77.64% done at 23626 kB/s, finished in 0:01
DUMP: Volume 1 completed at: Wed Sep 16 15:50:09 2015
DUMP: Volume 1 9132800 blocks (8918.75MB)
DUMP: Volume 1 took 0:07:01
DUMP: Volume 1 transfer rate: 21693 kB/s
DUMP: 9132800 blocks (8918.75MB)
DUMP: finished in 421 seconds, throughput 21693 kBytes/sec
DUMP: Date of this level 0 dump: Wed Sep 16 15:40:41 2015
DUMP: Date this dump completed: Wed Sep 16 15:50:09 2015
DUMP: Average transfer rate: 21693 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.

```
smw# Vi /mnt/etc/fstab
```

For example, change

```
/dev/disk/by-id/scsi-35000c50079ab34b7-part1 swap swap
defaults 0 0
```

```
/dev/disk/by-id/scsi-35000c50079ab34b7-part2 / ext3
acl,user_xattr 1 1
```

to:

```
/dev/disk/by-id/scsi-35000c50079ab71c4-part1 swap swap
defaults 0 0
/dev/disk/by-id/scsi-35000c50079ab71c4-part2 / ext3
acl,user_xattr 1 1
```

- 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/scsi-35000c50079ab34b7-part2
```

with:

```
root=/dev/disk/by-id/scsi-35000c50079ab71c4-part2
```

or 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` detects only the first 16 devices. Therefore, if a disk is added 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_1
```

- 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_1
find /THIS_IS_1
 (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
quit

```

**IMPORTANT:** For R630 SMWs, grub recreates `device.map` with the short names, not the persistent names. Do not trust the `/dev/sdx` names. Always use `find` when executing grub because it is possible that grub root may not be `hd2` the next time grub is executed.

#### 11. Unmount the backup root partition.

```
smw# umount /mnt
```

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

## R815 SMW: Create an SMW Bootable Backup Drive

### About this task

This procedure creates a bootable backup drive for a Dell R815 SMW in order to replace the primary drive if the primary drive fails. When these steps are completed, the backup drive on the SMW will be a bootable replacement for the primary drive when the backup drive is plugged in as or cabled as the primary drive.

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

**NOTICE:** To create a clean backup, Cray recommends shutting down the Cray system before beginning this procedure.

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

## Procedure

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

```

crayadm@smw> su -
Password:
smw#

```

#### 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 this site has completed this step at least once after installing the updated SMW base

operating system), skip to step 3 on page 121; 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 the R815 SMW.

- a. Identify the installed SMW drive model numbers, serial numbers, and associated Linux device (/dev) names.

Execute `smwmapdrives` 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 running that release or a later one, 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; these 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 desk-side 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
```

Cray recommends that `/boot/grub/device.map`, `/etc/fstab` and `/boot/grub/menu.lst` changes use the "by-path" rather than the "by-id" device name because that would allow physically swapping the backup drive into the primary slot when there is a disk failure in the primary disk. If the backup disk is intended as backup only, rather than as a bootable backup, it is acceptable to use either device name.

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

```
Dell Rackmount r815 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 step 2.a on page 119 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`.

For example, if the `root` and `swap` file systems are currently configured to mount `/dev/sda2`, they should be changed. Using the `by-path` device name from the example in step 2.a on page 119, the `fstab` lines would change from:

|                        |                   |                   |                             |                  |
|------------------------|-------------------|-------------------|-----------------------------|------------------|
| <code>/dev/sda1</code> | <code>swap</code> | <code>swap</code> | <code>defaults</code>       | <code>0 0</code> |
| <code>/dev/sda2</code> | <code>/</code>    | <code>ext3</code> | <code>acl,user_xattr</code> | <code>1 1</code> |



to:

```
/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
```

- 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-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 the 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 the `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-part2
```

or 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` detects only the first 16 devices. Therefore, if a disk is added 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 R815 SMWs, `grub` recreates `device.map` with the short names, not the persistent names. Do not trust the `/dev/sdx` names. Always use `find` when executing `grub` because it is possible that `grub root` may not be `hd2` the next time `grub` is executed.

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

### About this task

This procedure creates a System Management Workstation (SMW) bootable backup drive for a desk-side SMW. Its purpose is to replace the primary drive if the primary drive fails. When this procedure is completed, the backup drive on the SMW will be a bootable replacement for the primary drive when the backup drive is plugged in as or cabled as the primary drive.

**IMPORTANT:** The disk device names shown in this procedure are only examples; substitute the actual disk device names when executing this procedure. 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`.

**NOTICE:** To create a clean backup, Cray recommends shutting down the Cray system before beginning this procedure.

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

## Procedure

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 the 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. Be aware that errors in the `/boot/grub/menu.lst` will affect the 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 - 3.0.101-0.46
root (hd0,1)
kernel /boot/vmlinuz-3.0.101-0.46-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-3.0.101-0.46-default
###Don't change this comment - YaST2 identifier: Original name: failsafe###
title Failsafe -- SUSE Linux Enterprise Server 11 - 3.0.101-0.46
root (hd0,1)
kernel /boot/vmlinuz-3.0.101-0.46-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-3.0.101-0.46-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 the 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 - 3.0.101-0.46
kernel /boot/vmlinuz-3.0.101-0.46-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 - 3.0.101-0.46
kernel /boot/vmlinuz-3.0.101-0.46-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 - 3.0.101-0.46
kernel /boot/vmlinuz-3.0.101-0.46-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 - 3.0.101-0.46
kernel /boot/vmlinuz-3.0.101-0.46-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 - 3.0.101-0.46
root (hd0,1)
kernel /boot/vmlinuz-3.0.101-0.46-default root=/dev/sda2 \
resume=/dev/sda1 splash=silent crashkernel=256M-:128M@16M showopts
initrd /boot/initrd-3.0.101-0.46-default
###Don't change this comment - YaST2 identifier: Original name: failsafe###
title Failsafe -- SUSE Linux Enterprise Server 11 - 3.0.101-0.46
root (hd0,1)
kernel /boot/vmlinuz-3.0.101-0.46-default \
root=/dev/sda2 showopts ide=nodma apm=off noresume edd=off
powersaved=off nohz=off highres=off initrd /boot/initrd-3.0.101-0.46-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 a disk is added 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:** To update a system set that is running, shut down the Cray system before installing the update package.

## Shut Down the Cray System

### Procedure

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

### Procedure

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 the media type.

- To install the update package from disk, place the SMW 7.2UP04 Software DVD in the CD/DVD drive and mount it.

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

- 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.
- To install 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 being installed. For example:

```
smw# mount -o loop,ro smw-image-7.2.0-1.0702.37336.662-1.iso /media/cdrom
```

3. If `postfix` is not configured on the SMW, skip this step:

To prevent the `master.cf` and `main.cf` `postfix` configuration files from being recreated during software updates or fixes, ensure the following setting in the `/etc/sysconfig/mail` file on the 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 the new copy of the `SMWinstall.conf` file with site-specific requirements. Only `root` can 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 the installation 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 the installation 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 the installation site had power management profiles, copy the files that were 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 the installation 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 [Prepare to Upgrade or Update SMW Software](#) on page 109), compare the content of the local file to the distributed version of the file, and update the local file appropriately.
8. If the installation 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 the original `xtdiscover.ini` file, and update the `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 7.2UP04 Software media.

```
smw# umount /media/cdrom
```

10. If using the update disk, eject the SMW 7.2UP04 Software DVD.

```
smw# eject
```

## 11. Reboot the SMW.

```
smw# reboot
```

For a Cray XE or Cray XK system, continue to [Update the L0 and L1 Firmware](#) on page 138.

## For Cray XC Series Systems Only: Update the BC and CC Firmware

### About this task

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, go back to the last successful step in the installation procedure and continue from there. Do not skip steps or continue out of order.

### Procedure

#### 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 Mon Sep 14 00:57:20 CDT 2015 on hssbld0 by bwdev
lsb-cray-hss-7.2.0-1.0702.37336.662
smw# hss_make_default_initrd /opt/cray/hss-images/master/7.2.0-1.0702.37336.662
::: 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 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/
7.2.0-1.0702.37336.662.

Running hssclone.
Image Clone Complete: /opt/cray/hss-images/image-7.2.0-1.0702.37336.662
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-3.0.76-0.11.1_1.0702.8867-
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
```



- Exit from the `root` login.

```
smw# exit
```

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

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

If the system was not bounced previously, the following message may be displayed:

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

If so, respond with **y**.

- 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 *CLE Installation and Configuration Guide (S-2444)*.

- 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 there is uncertainty about 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 the `xtbounce --linktune` command is run.

```
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 System Hardware](#) on page 58.

## Update the L0 and L1 Firmware

### Prerequisites

This task applies to Cray XE and XK systems only.



**CAUTION:** Shut down (but do not power off) the Cray XE or XK system before beginning this task.

### Procedure

1. Log on as `crayadm`.
2. Flash the L1 and L0 controllers on the entire system (`s0`).

```
crayadm@smw> xtflash s0
```

Depending on the size of the system, the `xtflash` command may take a long time to run.

If all L0 daemons do not respond, ensure that the cable connections are secure or contact a Cray hardware representative for assistance.

3. As `root`, create a new boot image. For additional information, see *Installing and Configuring Cray Linux Environment (CLE) Software* (S-2444), which is provided with the CLE release package. The name of the script that is used to build boot images is `/var/opt/cray/install/shell_bootimage_LABEL.sh`, where `LABEL` is the system set label. The system set label can follow any naming convention, but typically it will identify the version number and/or the date the image was created.

```
crayadm@smw> su - root
smw# /var/opt/cray/install/shell_bootimage_label_name.sh
```

4. Exit from `root`, and bounce the L1 and L0 controllers on the entire system (`s0`).

```
smw# exit
crayadm@smw> xtbounce s0
```

## For Cray XE or Cray XK Systems Only: Verify Hardware Programmable Interface Controllers (PIC)

### About this task

**RESTRICTION:** This task applies only to Cray XE and Cray XK systems.

### Procedure

Execute the `xtcheckpic` command to verify that the PICs are current.

If any PICs need to be updated, contact 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, contact 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 the installation site does not install the latest CLE release package.

## Update SMW Software on the Boot Root and Shared Root

### About this task

This procedure uses the `SMWinstallCLE` script to update the SMW software on the boot root and shared root for systems already running the Cray Linux Environment (CLE) software. The RPMs that `SMWinstallCLE` installs on the boot root and shared root will also be installed when `CLEinstall` runs during a CLE update.

**TIP:** Use this procedure only if the plan is to boot CLE after the SMW update but before updating the CLE software. Otherwise, if the plan is to update CLE software without booting CLE after the SMW update, it is safe to skip this procedure.

For more information about the `SMWinstallCLE` script, see the `SMWinstallCLE(8)` man page.

### Procedure

1. As `root`, mount the release media by using one of the following commands, depending on the media type.

- To install the update package from disk, place the SMW 7.2UP04 Software DVD in the CD/DVD drive and mount it.

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

- 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.
- To install 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 being installed. For example:

```
smw# mount -o loop,ro smw-image-7.2.0-1.0702.37336.662-1.iso /media/cdrom
```

2. Update the `label_name` system set from the `/etc/sysset.conf` system set configuration file.

In the following steps it is assumed that the label `label_name` is described in the `/etc/sysset.conf` system set configuration file. See the `sysset.conf(5)` man page for additional information about the `/etc/sysset.conf` file. For more detailed information about `SMWinstallCLE`, see the `SMWinstallCLE(8)` man page.

```
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 for confirmation that the system set being changed 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
```

3. Unmount the SMW 7.2UP04 Software media.

```
smw# umount /media/cdrom
```

4. If using the update disk, eject the SMW 7.2UP04 Software DVD.

```
smw# eject
```

The SMW software is now updated. If the firewall is not yet configured, see [Set Up the SUSE Firewall and IP Tables](#) on page 61. Then continue to install the CLE software using *CLE Installation and Configuration Guide* (S-2444), which is provided with the CLE release package.

**NOTE:** To reconfigure a LOGDISK, PMDISK, or DBDISK when it is necessary to replace a failed drive with a new drive on a rack-mount SMW, see [Replace a Failed Disk Drive](#) on page 153.

## Set Up the SUSE Firewall and IP Tables

### About this task

The SMW software includes a firewall. The following steps enable and configure the firewall.

**TIP:** It is not necessary to shut down the system before performing this task.

### Procedure

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. Using the SuSEfirewall12 program and the following steps, 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 10.3.1.1,udp,514 10.3.1.254,tcp,
6811:6815"
```

**NOTE:** 10.3.1.254 is the boot node's IP address for eth0 on the network between the boot node and the SMW.

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. Verify the 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 using VNC](#) on page 147.

## Configure the Simple Event Correlator (SEC)

The System Management Workstation (SMW) 7.2.UP04 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.

To use the Cray SEC, see *Configure SEC Software* (S-2542) for configuration procedures.

# Perform Additional SMW Installation or Configuration Procedures

---

The following procedures may be performed, as needed, after the essential SMW installation or update/upgrade procedures are complete.

- [Install Additional Software as RPMs](#) on page 143: How to install third-party RPM packages.
- [Change the Default iDRAC Password](#) on page 60: How to change the default password for an Integrated Dell Remote Access Controller (iDRAC) on an R630 or R815 rack-mount SMW.
- [Use the iDRAC](#) on page 143: How to use the iDRAC after it has been enabled on an R630 or R815 rack-mount SMW.
- [Change the Time Zone for the SMW and the Blade and Cabinet Controllers](#) on page 144: How to change the time zone for the SMW and the blade and cabinet controllers in Cray XC30, XE, and XK systems. Includes instructions for first-time and subsequent time zone modification.
- [Rack-mount SMW: Install the Dell Systems Management Tools and Documentation DVD](#) on page 147: How to enable access to advanced features such as automated recovery by installing Dell management tools and documentation on a rack-mount SMW.

## 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.

## Use the iDRAC

### About this task

#### Procedure

1. Bring up a web browser.
2. Go to: `https://cray-drac`. A login screen appears.
3. Enter the account user name and password set up in , [R630 SMW: Change the BIOS and iDRAC Settings](#) on page 16, or [Change the Default iDRAC Password](#) on page 60.

The **System Summary** window appears.

4. Select **Submit**.

5. To access the SMW console, select the **Console Media** tab.

The **Virtual Console and Virtual Media** window appears.

6. Select **Launch Virtual Console**.

**TIP:** By default, the console window has two cursors: one for the console and one for the administrator's 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.

**TIP:** To logout of the virtual console, kill the window or select **File**, then **Exit**. The web browser is still logged into the iDRAC.

For detailed information, download the iDRAC documentation at: <http://www.dell.com/support>.

## Change the Time Zone for the SMW and the Blade and Cabinet Controllers

### About this task



**CAUTION:** Perform this procedure only when the Cray system is shut down. Do not flash blade and cabinet controllers while the Cray system is booted.

In this example, the time zone is changed from "America/Chicago" to "America/New\_York".

### Procedure

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

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

2. Ensure the blade and cabinet controllers are responding. For example:

```
smw# xtalive -a 10sysd s0
```

3. Check the current time zone setting for the SMW and controllers.

```
smw# date
Wed Aug 01 21:30:06 CDT 2012
smw# xtrsh -l root -s /bin/date s0
c0-0c0s2 : Wed Aug 01 21:30:51 CDT 2012
c0-0c0s5 : Wed Aug 01 21:30:51 CDT 2012
c0-0c0s7 : Wed Aug 01 21:30:51 CDT 2012
c0-0c1s1 : Wed Aug 01 21:30:51 CDT 2012
.
.
.
c0-0 : Wed Aug 01 21:30:52 CDT 2012
```



4. Verify that the `zone.tab` file in the `/usr/share/zoneinfo` directory contains the desired time zone.

```
smw# grep America/New_York /usr/share/zoneinfo/zone.tab
US +404251-0740023 America/New_York Eastern Time
```

5. Create the time conversion information files.

```
smw# date
Wed Aug 01 21:32:52 CDT 2012
smw# /usr/sbin/zic -l America/New_York
smw# date
Wed Aug 01 22:33:05 EDT 2012
```

6. Modify the `clock` file in the `/etc/sysconfig` directory to set the `DEFAULT_TIMEZONE` and the `TIMEZONE` variables to the new time zone.

```
smw# grep TIMEZONE /etc/sysconfig/clock
TIMEZONE="America/Chicago"
DEFAULT_TIMEZONE="US/Eastern"
smw# vi /etc/sysconfig/clock
make changes
smw# grep TIMEZONE /etc/sysconfig/clock
TIMEZONE="America/New_York"
```

7. Copy the `/etc/localtime` file to `/opt/tftpboot`, and then restart the log system and `rsms`.

```
smw# cp /etc/localtime /opt/tftpboot
smw# /etc/init.d/cray-syslog restart
smw# /etc/init.d/rsms restart
```

For Cray XC30 systems, continue with step [10](#) on page 146.

8. **NOTE:** For Cray XE or Cray XK systems only:

If this is the first time the time zone has been modified, complete this step. If the time zone has been changed already, skip this step and perform step [9](#) on page 146.

- a. Exit from the `root` login.

```
smw# exit
```

- b. Erase the flash memory of the L1s and flash the updated time zone.

```
crayadm@smw> fm -w -t 11
crayadm@smw> xtflash -t 11
```

- c. Erase the flash memory of the L0s and flash the updated time zone.

```
crayadm@smw> fm -w -t 10
crayadm@smw> xtflash -t 10
```

- d. Check the current time zone setting for the SMW and controllers.

```
crayadm@smw> date
Wed Aug 01 23:07:07 EDT 2012
crayadm@smw> xtrsh -l root -s /bin/date s0
c0-0c1s1 : Wed Aug 01 23:07:16 EDT 2012
c0-0c0s7 : Wed Aug 01 23:07:16 EDT 2012
c0-0c1s3 : Wed Aug 01 23:07:16 EDT 2012
```

```
.
.
.
c0-0 : Wed Aug 01 23:07:17 EDT 2012
```

Continue with step [11](#) on page 146.

**9. NOTE:** For Cray XE or Cray XK systems only:

If the time zone has been changed already, complete this step. If this is the first time the time zone has been modified, perform step [8](#) on page 145.

- a. To update the L1's time zone:

```
smw# xtrsh -l root -m ^c[0-9]+-[0-9]+$ -s 'atftp -g -r localtime \
-l $(readlink /etc/localtime) router && cp /etc/localtime /var/tftp'
```

- b. To update the L0's time zone:

```
smw# xtrsh -l root -m s -s 'atftp -g -r localtime \
-l $(readlink /etc/localtime) router'
```

Continue with step [11](#) on page 146.

**10. NOTE:** For Cray XC30 systems only:

Reboot the cabinet controllers to get the updated time zone.

- a. Power down the system.

```
smw# xtcli power down s0
```

- b. 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
```

- c. 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.
```

- d. Exit from the root login.

```
smw# exit
```

**11. Bounce the system.**

```
crayadm@smw> xtbounce s0
```

**NOTE:** An incompatibility exists between the current version of `/etc/localtime` and earlier versions that may be on the system. This incompatibility causes the `date` command to report an incorrect time on the compute nodes. To resolve this incompatibility, after updating the SMW software, update the time zone on the compute nodes also as described in the procedure *Changing the time zone for compute nodes* in *CLE Installation and Configuration Guide*.

## Rack-mount SMW: Install the Dell Systems Management Tools and Documentation DVD

### About this task

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 OpenManage Server Administrator software from the Dell Systems Management Tools and Documentation DVD. For rack-mount SMWs, this DVD is provided with the system.

Visit the Dell OpenManage Linux Repository to view the Dell OpenManage Server Administrator documentation: <http://linux.dell.com/wiki/index.php/Repository/OMSA>

### Procedure

1. Obtain the Dell System Management Tools and Documentation DVD.
2. Log on to the SMW as `root`.
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.

The system can now be managed for Properties, Shutdown, Logs, Alert Management, and Session Management.

## Enable Remote Access to the SMW using VNC

Virtual network computing (VNC) software enables a user to view and interact with the SMW from another computer. The Cray system provides a VNC server, `Xvnc`, and a VNC account, `cray-vnc`. Enabling VNC is a site choice. It is more useful for the desk-side SMWs than for the rack-mount SMWs because the rack-mount SMWs have the iDRAC, which can be used for remote interaction with the full graphical console. Cray recommends this procedure for desk-side SMWs but not for rack-mount SMWs.

- To start the VNC server, see [Start the VNC server](#) on page 148.
- To obtain a VNC client to connect to the server, download a VNC client from a reputable website, such as those listed here:
  - RealVNC™: <http://www.realvnc.com/>
  - TightVNC: <http://www.tightvnc.com/>

**ATTENTION:** 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 Cray recommends that the VNC client only connect to the VNC server on the SMW via an SSH tunnel.

To create a secure VNC connection, see these topics on connecting to the VNC server through an SSH tunnel:

- [For Workstation or Laptop Running Linux: Connecting to the VNC Server through an SSH Tunnel, using the vncviewer -via Option](#) on page 149
- [For Workstation or Laptop Running Linux: Connecting to the VNC Server through an SSH Tunnel](#) on page 150
- [For Workstation or Laptop Running Mac OS X: Connecting to the VNC Server through an SSH Tunnel](#) on page 150
- [For Workstation or Laptop Running Windows: Connecting to the VNC Server through an SSH Tunnel](#) on page 151

## Start the VNC server

### Procedure

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 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 a `failed` notification, this is normal and can be ignored.

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 for one. 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`. With many clients, it is possible to specify whether to connect in view-only or in an active mode. If active participation is specified, every mouse movement and keystroke made in the client is sent to the server. If more than one client is active at the same time, 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

### About this task

**IMPORTANT:** This procedure is for use with the TightVNC client program.

Verify that the `vncviewer -via` option is available. If not, use [For Workstation or Laptop Running Linux: Connecting to the VNC Server through an SSH Tunnel](#) on page 150.

### Procedure

To connect from a workstation or laptop running Linux, enter the `vncviewer` command shown below.

The first password to enter is for `crayadm` on the SMW. The second password to enter is for the VNC server on the SMW, which was set the 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

### About this task

**NOTE:** This procedure assumes that the VNC server on the SMW is running with the default port of 5901.

### Procedure

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. The system will prompt 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. The system will prompt 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

### About this task

**NOTE:** This procedure assumes that the VNC server on the SMW is running with the default port of 5901.

### Procedure

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. The system will prompt 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. The system will prompt 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

### About this task

**NOTE:** To connect from a computer running Windows, both a VNC client program, such as TightVNC, and an SSH program, such as PuTTY, SecureCRT, or OpenSSH are recommended.

### Procedure

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.

## Troubleshoot SMW Problems

The following procedures address issues that may occur during installation or operation of a Cray SMW.

- [R630 SMW: System Repair](#) on page 152: How to reach the menu options for repairing or rescuing an installed system on a Dell R630 SMW.
- [R815 SMW: Configure the SMW to Restore AC Power After a Power Failure](#) on page 152: How to configure the SMW to restore AC power whenever a power failure occurs.
- [Replace a Failed Disk Drive](#) on page 153: How to replace a failed LOGDISK, DBDISK, or PMDISK on a rack-mount SMW.
- [Recover from a Corrupt or Missing HSS MySQL Database](#) on page 157: How to restore and/or regenerate a missing or damaged MySQL database if attempts by MySQL to repair the tables fail.
- [Troubleshoot Temperature Warnings Reported in an End Cabinet](#) on page 159: How a small entry in the hss.ini file can solve end-cabinet temperature problems in liquid-cooled systems. This procedure can be performed proactively to avoid cooling problems.

### R630 SMW: System Repair

If there is occasion to repair or rescue a system installed on an R630 SMW, it is necessary to first boot from the Dell driver update disk (DUD) DVD, select the appropriate menu option (shown below), then insert the Cray-SMWbase11SP3-201507081500 DVD when prompted. If the SMW is booted from the Cray-SMWbase11SP3 DVD first, the appropriate menu options do not appear.

```
Repair Installed System
Rescue Installed System
```

By contrast, the following menu options appear regardless of whether the Dell DUD DVD or the Cray-SMWbase11SP3 DVD was used to boot the system. These options are useful for checking the system.

```
Firmware Test
Memory Test
```

### R815 SMW: Configure the SMW to Restore AC Power After a Power Failure

#### About this task

#### Procedure

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.

## Replace a Failed Disk Drive

The procedures that follow describe how to replace a failed LOGDISK, DBDISK, or PMDISK. These procedures require 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.

- For a Dell R630 rack-mount SMW:
  - [R630 SMW: Replace a Failed LOGDISK](#) on page 153
  - The procedure to replace a failed DBDISK on a Dell R630 SMW is currently in development. This publication will be revised to include it as soon as possible.
  - The procedure to replace a failed PMDISK on a Dell R630 SMW is currently in development. This publication will be revised to include it as soon as possible.
- For a Dell R815 rack-mount SMW:
  - [R815 SMW: Replace a Failed LOGDISK](#) on page 155
  - The procedure to replace a failed DBDISK on a Dell R815 SMW is currently in development. This publication will be revised to include it as soon as possible.
  - The procedure to replace a failed PMDISK on a Dell R815 SMW is currently in development. This publication will be revised to include it as soon as possible.

## R630 SMW: Replace a Failed LOGDISK

### About this task

This procedure describes how to replace a failed LOGDISK on a Dell R630 SMW. To replace a failed LOGDISK on a Dell R815 SMW, see [R815 SMW: Replace a Failed LOGDISK](#) on page 155.

## Procedure

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-0000:03:00.0-scsi-0:0:3:0
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-0000:03:00.0-scsi-0:0:3:0:1000.2 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-0000:03:00.0-scsi-0:0:3:0-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-0000:03:00.0-scsi-0:0:3:0-part1
```

5. Mount the newly created file system.

```
smw# mount /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:3:0-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
```

## R815 SMW: Replace a Failed LOGDISK

### About this task

This procedure describes how to replace a failed `LOGDISK` on a Dell R815 SMW. To replace a failed `LOGDISK` on a Dell R630 SMW, see [R630 SMW: Replace a Failed LOGDISK](#) on page 153.

### Procedure

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

```

## Recover from a Corrupt or Missing HSS MySQL Database

If the HSS MySQL database has been damaged or is missing, there are several possible courses of action.

- **Repair.** If the database has become corrupt, MySQL automatically attempts to repair damaged tables. Look in the log file (default `/var/lib/mysql/machine.err`) for suggested manual recovery steps, if any, and try those first. Repairing the database is the best option when possible.
- **Restore and regenerate.** If there are no suggestions or the suggested steps fail to repair the database, use the procedure [Restore a Corrupt or Missing HSS MySQL Database from a Backup](#) on page 157. Restoring the database from the most recent backup (provided a more recent manual backup is not available) will restore the database to its state just prior to the last `xtdiscover` or `warmswap` add operation. An incremental discovery to the present system state will usually be faster than one made from a fresh database, and it will not require administrative state changes made prior to the backup (such as marking compute nodes as 'service') to be performed again.

**TIP:** To minimize needed discovery, make more frequent backups:

```
/usr/bin/mysqldump --add-drop-database --routines -uhssds -phssds hssds
> /home/crayadm/hss_db_backup/my-new-hssds-backup.sql
```

The HSS MySQL database could be backed up after every successful `warmswap` (`xtdiscover --warmswap`), regular `xtdiscover`, and any administrative state change (e.g., `xtcli disable/enable/set_empty/mark_node`). Because these actions are all logged in the commands log, they could be used to automatically trigger backups.

- **Regenerate from scratch.** If all else fails, use the procedure [Regenerate a Corrupt or Missing HSS MySQL Database from Scratch](#) on page 158. In this case, the database and the database root password are wiped out, and discovery is used to regenerate the database.

## Restore a Corrupt or Missing HSS MySQL Database from a Backup

### About this task

If the HSS MySQL database becomes corrupt or is missing, and automated attempts to repair damaged tables have failed, use this procedure to do a partial restoration from backup.

### Procedure

1. Stop the HSS daemons (by stopping RSMS) and the MySQL service.

```
smw> rsms stop
smw> cd /var/lib/mysql
smw> sudo /etc/init.d/mysql stop
```

2. Move the damaged database files out of the MySQL directory.

```
smw> mkdir /tmp/backup11
smw> sudo mv ibdata1 ib_logfile0 ib_logfile1 hssds /tmp/backup11
```

This procedure assumes that the old database files cannot be repaired; however, this step retains those old database files (just in case) and clears out the MySQL directory.

3. Restart MySQL.

```
smw> sudo /etc/init.d/mysql start
```

4. Ensure the database is gone.

```
smw> mysql -uhssds -phssds -e "drop database hssds"
```

If the database is gone, the following error message appears:

```
ERROR 1008 (HY000) at line 1: Can't drop database 'hssds'; database doesn't exist
```

5. Load the most recent backup (from `/home/crayadm/hss_db_backup/`).

```
smw> mysql -uhssds -phssds < db_backup.11-17-2014.1120.sql
```

The backups in `/home/crayadm/hss_db_backup/` are from past runs of `xtdiscover` and `xtwarmswap --add` and were taken *before* the state of the database was updated.

6. Restart the HSS daemons (important!) and then use `xtdiscover` to pick up any changes to the system since the backup was taken (or all of the database, if a backup was not loaded in the previous step).

```
smw> rsms start
smw> sudo xtdiscover
```

## Regenerate a Corrupt or Missing HSS MySQL Database from Scratch

### About this task

If the HSS MySQL database becomes corrupt or is missing, and all attempts to repair or restore it have failed, use this procedure to regenerate the database from scratch. Deleting `/var/lib/mysql` removes everything that stores MySQL state, including the password. Hence the need to recreate it in step x. When MySQL is restarted and the MySQL directory does not exist, `/var/lib/mysql` will be re-created as new.

### Procedure

1. Stop the HSS daemons (by stopping RSMS) and the MySQL service.

```
smw> rsms stop
smw> sudo /etc/init.d/mysql stop
```

2. Remove the damaged database.

```
smv> sudo mv /var/lib/mysql /var/lib/mysql.bad
```

Remove the MySQL directory by renaming it. A newly initialized `/var/lib/mysql` directory will be created when the MySQL service is restarted.

3. Restart MySQL.

```
smw> sudo /etc/init.d/mysql start
```

The MySQL directory is reinitialized, and the default password is set to the empty string.

4. Reset the MySQL root password to its former value.

```
smw> mysqladmin -uroot password -p
```

The system prompts for the existing password. Press the **Enter** key to enter an empty string (the default password set at reinitialization). When the **New password** prompt appears, enter the former password (the root password before the system was reinitialized), and enter it again at the **Confirm new password** prompt.

5. Initialize the HSS database tables and restore user permission tables.

```
smw> hssds_init
smw> dbgrant
```

The system will prompt for a password after each of the above two commands. Give the newly reset MySQL root password each time.

6. Restart the HSS daemons (important!) and then discover to regenerate the database.

```
smw> rsms start
smw> sudo xtdiscover --bootstrap
smw> sudo xtdiscover
```

## Troubleshoot Temperature Warnings Reported in an End Cabinet

### About this task

If the consumer log or `xtcheckhss` reports temperature warnings in an end-of-row cabinet of a liquid-cooled system, the current `hss.ini` file may not have the necessary temperature set point defined, or the set point value may not be appropriate for the site. For more information, see [About End-Cabinet Cooling](#) on page 160.

### Procedure

If the consumer log shows entries similar to the example below, perform the procedure in [Ensure End-Cabinet Temperature Set Point is Defined](#) on page 160.

```
Mon Jul 28 05:59:47 2014 - rs_event_t at 0x7f5bc0000920
ev_id = 0x080040ed (ec_ll_failed)
ev_src = ::c1-0
ev_gen = ::c0-0c0s0n0
ev_flag = 0x00000002 ev_priority = 0 ev_len = 158 ev_seqnum = 0x00000000
ev_stp = 53d5e6d3.0000176d [Mon Jul 28 05:59:47 2014]
svcid 0: ::c1-0 = svid_inst=0x0/svid_type=0x0/svid_node=c1-0[rsn_node=0x0/
rsn_type=0x3/rsn_state=0x6], err code 65914
- Cabinet Controller Temperature Fault
ev_data...
00000000: 01 00 00 00 00 00 00 00 00 00 00 00 00 00 0c 06 00 00 *.....*
00000010: 04 00 00 00 00 00 00 00 00 00 01 00 00 00 7a 01 01 00 *.....z...*
00000020: 7a 00 00 00 30 39 34 7c 57 41 52 4e 7c 54 45 4d *z...094|WARN|TEM*
00000030: 50 7c 2f 64 61 74 61 2f 63 6f 6d 70 75 74 65 5f *P|/data/compute_*
00000040: 63 61 62 69 6e 65 74 2f 61 69 72 5f 73 65 6e 73 *cabinet/air_sens*
00000050: 6f 72 73 2f 63 68 32 2f 61 69 72 5f 74 65 6d 70 *ors/ch2/air_temp*
00000060: 32 3a 64 65 67 63 2a 31 30 30 7c 4d 61 78 69 6d *2:degc*100|Maxim*
00000070: 75 6d 20 73 6f 66 74 20 6c 69 6d 69 74 20 65 78 *um soft limit ex*
00000080: 63 65 65 64 65 64 21 7c 44 61 74 61 3d 33 30 30 *ceeded!|Data=300*
00000090: 32 7c 4c 69 6d 69 74 3d 33 30 30 30 2e 00 *2|Limit=3000....*
```

With `xtcheckhss`, this problem looks like:

```
No Version Mismatches Found!
=====
===== Sensor Warnings =====
=====
Component Module Sensor HMIN SMIN DATA UNIT SMAX HMAX

c2-0 compute_cabinet ambient_temp0 30 50 324 degc*10 300 350
c2-0 compute_cabinet ambient_temp1 30 50 306 degc*10 300 350
c2-0 compute_cabinet ch0_air_temp0 0 1000 3486 degc*100 3000 3500
c2-0 compute_cabinet ch0_air_temp1 0 1000 3355 degc*100 3000 3500
c2-0 compute_cabinet ch0_air_temp2 0 1000 3338 degc*100 3000 3500
c2-0 compute_cabinet ch0_air_temp3 0 1000 3486 degc*100 3000 3500

No SEEP Errors Found!
No ITP Errors Found!
No NTP Time Sync Errors Found!
No Control Errors Found!
```

## About End-Cabinet Cooling

In a liquid-cooled cabinet with chassis (cages) that are unevenly populated, the exit temperatures in each cage will be very different. In a normal cabinet, the water valve is controlled by the average temperature of the hottest temperature strip. By contrast, the water valve in an end-of-row cabinet is controlled by the average temperature of all temperature strips. This may lead to inadequate cooling of a populated cage if the other two cages are not populated or have minimal heat load.

To avoid problems arising from inadequate cooling, the exit air temperatures of the end-of-row cabinet can be independently controlled. This is achieved through an entry in the `hss.ini` file that sets the end-of-row cabinet exit temperature lower than that of other cabinets. The default value is 22°C; however this should be adjusted to meet site-specific requirements. If the end cabinet exit air temperature is not defined in the `hss.ini` file, the air temperature will default to the setting defined for the other cabinets in the cooling row.

## Ensure End-Cabinet Temperature Set Point is Defined

### About this task

Cray added a temperature set point for end cabinets to the `hss.ini` file in the SMW 7.2.UP01 release for liquid-cooled systems. Use this procedure to ensure that this temperature set point is defined in the current `hss.ini` and is set to an appropriate value.

## Procedure

1. Open the `hss.ini` and look for the following entry. If it is not present, add it.

```
This group is used to define the attributes that are only applied to the
end cabinet of a row. The attributes defined here will override the same
attributes in group [ccrd] above. If no attributes are defined in this group
the end cabinet will be configured using the attributes of group [ccrd].
[endcabinet]
Define temperature setpoints. This controls the temperature of the air that is
being exhausted into the room.
18C to 27C - Change it to match site requirement.
temp_setpoint=22
```



2. Adjust the value of `temp_setpoint` as appropriate for the installation site.

To determine an appropriate value, consider the following:

- the inlet water temperature, which should be below the exit air temperature setting
- the facility room environment