



Installing Cray System Management Workstation (SMW) Software

S-2480-7001

© 2007–2012 Cray Inc. All Rights Reserved. This document or parts thereof may not be reproduced in any form unless permitted by contract or by written permission of Cray Inc.

U.S. GOVERNMENT RESTRICTED RIGHTS NOTICE

The Computer Software is delivered as "Commercial Computer Software" as defined in DFARS 48 CFR 252.227-7014.

All Computer Software and Computer Software Documentation acquired by or for the U.S. Government is provided with Restricted Rights. Use, duplication or disclosure by the U.S. Government is subject to the restrictions described in FAR 48 CFR 52.227-14 or DFARS 48 CFR 252.227-7014, as applicable.

Technical Data acquired by or for the U.S. Government, if any, is provided with Limited Rights. Use, duplication or disclosure by the U.S. Government is subject to the restrictions described in FAR 48 CFR 52.227-14 or DFARS 48 CFR 252.227-7013, as applicable.

Cray and Sonexion are federally registered trademarks and Active Manager, Cascade, Cray Apprentice2, Cray Apprentice2 Desktop, Cray C++ Compiling System, Cray CX, Cray CX1, Cray CX1-iWS, Cray CX1-LC, Cray CX1000, Cray CX1000-C, Cray CX1000-G, Cray CX1000-S, Cray CX1000-SC, Cray CX1000-SM, Cray CX1000-HN, Cray Fortran Compiler, Cray Linux Environment, Cray SHMEM, Cray X1, Cray X1E, Cray X2, Cray XD1, Cray XE, Cray XEm, Cray XE5, Cray XE5m, Cray XE6, Cray XE6m, Cray XK6, Cray XK6m, Cray XMT, Cray XR1, Cray XT, Cray XTm, Cray XT3, Cray XT4, Cray XT5, Cray XT5_h, Cray XT5m, Cray XT6, Cray XT6m, CrayDoc, CrayPort, CRInform, ECOphlex, LibSci, NodeKARE, RapidArray, The Way to Better Science, Threadstorm, uRiKA, UNICOS/lc, and YarcData are trademarks of Cray Inc.

Aries, Gemini, and SeaStar are trademarks of Intel Corporation in the United States and/or other countries. Dell, OpenManage, and PowerEdge are trademarks of Dell, Inc. DDN and S2A are trademarks of DataDirect Networks, Inc. Engenio, NetApp, and SANtricity are trademarks of NetApp, Inc. in the United States and/or other countries. Flash is a trademark of Adobe Inc. GNU is a trademark of The Free Software Foundation. ISO is a trademark of International Organization for Standardization (Organisation Internationale de Normalisation). Java, Lustre, MySQL, and Solaris are trademarks of Oracle and/or its affiliates. Linux is a trademark of Linus Torvalds. LSI is a trademark of LSI Corporation. Mac, Mac OS, and OS X are trademarks of Apple, Inc. QLogic, and SANbox are trademarks of QLogic Corporation. SLES and SUSE are trademarks of Novell, Inc. VM is a trademark of International Business Machines Corporation. Windows is a trademark of Microsoft Corporation. UNIX, the "X device," X Window System, and X/Open are trademarks of The Open Group in the United States and other countries. All other trademarks are the property of their respective owners.

RECORD OF REVISION

S-2480-7001 Published November 2012 Supports the Cray System Management Workstation (SMW) 7.0.UP01 release. The SMW 7.0.UP01 release also supports the Limited Availability (LA) release for Cray Cascade systems.

S-2480-70 Published August 2012 Supports the Cray System Management Workstation (SMW) 7.0.UP00 Limited Availability (LA) release.

S-2480-6002 Published December 2011 Supports the Cray System Management Workstation (SMW) 6.0.UP02 release.

S-2480-60 Published June 2011 Supports the Cray System Management Workstation (SMW) 6.0 release.

S-2480-5102 Published December 2010 Supports the Cray System Management Workstation (SMW) 5.1.UP02 release; it also provides procedures to upgrade the SMW base operating system from SUSE Linux Enterprise Server 11 (SLES 11) to SLES 11 Service Pack 1 (SLES 11 SP1).

S-2480-5101 Published September 2010 Supports the Cray System Management Workstation (SMW) 5.1.UP01 release.

5.1 Published June 2010 Supports the Cray System Management Workstation (SMW) 5.1 release.

5.0 Published March 2010 Supports the Cray System Management Workstation (SMW) 5.0 release.

4.0 Published May 2009 Supports general availability (GA) release of the Cray System Management Workstation (SMW) 4.0 release.

4.0 Published February 2009 Supports limited availability (LA) release of the Cray System Management Workstation (SMW) 4.0 release.

3.1 Published July 2008 Supports general availability (GA) release of the Cray System Management Workstation (SMW) 3.1 release.

3.1 Published December 2007 Supports limited availability (LA) release of the Cray System Management Workstation (SMW) 3.1 release.

3.0.1 Published October 2007 Supports general availability (GA) release of the Cray System Management Workstation (SMW) 3.0.1 release.

Contents

	<i>Page</i>
Introduction [1]	11
1.1 Related Publications	11
1.2 Distribution Media	12
Installing the SMW Release Package [2]	13
2.1 Preparing to Install the SMW Release Software	13
2.1.1 Before You Start the SMW Release Installation	13
2.1.1.1 Network Connections	14
2.1.1.2 Configuration Values	14
2.1.1.3 Passwords	16
2.1.2 Shutting Down the Cray System	16
2.2 Rack-mount SMW: Installing the SMW Base Operating System	17
2.3 Desk-side SMW: Installing the SMW Base Operating System	25
2.4 Installing and Configuring the SMW Software Packages	30
2.5 Discovering the Cray System Hardware and Powering Up the Full System	33
2.6 Confirming the SMW is Communicating with the System Hardware	46
2.7 For Cray Systems with Cray XK Blades: Updating the Cray XK Accelerator Firmware	47
2.8 Changing Default SMW Passwords After Completing Installation	48
2.9 Setting Up the SUSE Firewall and IP Tables	48
2.10 Collecting Additional Data During a System Dump Session	48
2.11 Backing Up Your Newly-installed SMW Software	48
Configuring Boot RAID [3]	49
3.1 Prerequisites and Assumptions for Configuring the Boot RAID	49
3.2 Configuring and Zoning the Boot RAID Using a DDN Storage System	50
3.2.1 Using the Command-line Interface (telnet) to Configure a DDN Storage System for Your Boot RAID	50
3.2.1.1 Configuring the Boot RAID Password When Using DDN Devices	50
3.2.1.2 Configuring the LUNs for DDN Devices	51
3.2.1.3 Zoning the LUNs for DDN Devices	53

	<i>Page</i>
3.2.2 Installing the Silicon Storage Appliance Manager GUI to Configure a DDN Storage System for Your Boot RAID	56
3.2.2.1 Installing the DDN Silicon Storage Appliance Manager	56
3.2.2.2 Configuring and Zoning the LUNs for DDN Devices	58
3.2.3 Configuring Remote Logging of DDN Messages	58
3.3 Configuring the Boot RAID Using a NetApp, Inc. Engenio Storage System	59
3.3.1 Installing SANtricity Storage Manager Software for NetApp, Inc. Engenio Devices	59
3.3.2 Configuring the LUNs for NetApp, Inc. Engenio Devices	61
3.3.3 Configuring Remote Logging of NetApp, Inc. Engenio Storage System Boot RAID Messages	63
3.4 Zoning the QLogic FC Switch	63
3.5 Rediscovering the LUNs	66
3.6 Partitioning the LUNs	66
Updating or Upgrading Your Cray SMW Software [4]	67
4.1 Before You Start	67
4.2 Backing Up Your Current Software	67
4.3 Installing an SMW Update Package	68
4.4 For Cray Cascade Systems Only: Updating the BC and CC Firmware	71
4.5 For Cray XE or Cray XK Systems Only: Updating the L0 and L1 Firmware	74
4.6 Confirming the SMW is Communicating with the System Hardware	75
4.7 For Cray XE or Cray XK Systems Only: Verifying Hardware Programmable Interface Controllers (PICs)	76
4.8 For Cray Systems with Cray XK Blades Only: Updating the Cray XK Accelerator Firmware	77
4.9 Updating SMW Software on the Boot Root and Shared Root	77
Appendix A Installing RPMs	79
A.1 Generic RPM Usage	79
Appendix B Updating the Time Zone	81
Appendix C Remote Access to the SMW	85
Appendix D Setting Up the SUSE Firewall and IP Tables	89
Appendix E Rack-mount SMW: Creating a Bootable Backup Drive	91
E.1 Rack-mount SMW: Creating a Cray System Management Workstation (SMW) Bootable Backup Drive	91
Appendix F Rack-mount SMW: Setting up the Bootable Backup Drive as an Alternate Boot Device	103
F.1 Rack-mount SMW: Setting Up the Bootable Backup Drive as an Alternate Boot Device	103

	<i>Page</i>
Appendix G Desk-side SMW: Creating a Bootable Backup Drive	107
G.1 Desk-side SMW: Creating an System Management Workstation (SMW) Bootable Backup Drive . . .	107
Appendix H Desk-side SMW: Setting Up the Bootable Backup Drive as an Alternate Boot Device	117
H.1 Desk-side SMW: Setting Up the Bootable Backup Drive as an Alternate Boot Device	117
Appendix I Enabling an Integrated Dell Remote Access Controller (iDRAC6) on a Rack-mount SMW	121
I.1 Before You Start	121
I.2 Enabling an Integrated Dell Remote Access Controller (iDRAC6) on a Rack-mount SMW	121
I.3 Using the iDRAC6	122
Appendix J Rack-mount SMW: Installing Server Administrator Software	125
Appendix K Rack-mount SMW: Replacing a Failed LOGDISK or DBDISK Disk Drive	127
K.1 Rack-mount SMW: Replacing a Failed LOGDISK or DBDISK Disk Drive	127
Procedures	
Procedure 1. Shutting down the Cray system	16
Procedure 2. Rack-mount SMW: Installing the base operating system	17
Procedure 3. Desk-side SMW: Installing the base operating system	25
Procedure 4. Installing and configuring SMW software	30
Procedure 5. Bootstrapping hardware discovery	33
Procedure 6. Discovering the Cray system hardware	39
Procedure 7. Confirming the SMW is communicating with the system hardware	46
Procedure 8. Configuring the boot RAID password when using DDN devices	50
Procedure 9. Configuring the LUNs for DDN devices	51
Procedure 10. Zoning the LUNs for DDN devices	53
Procedure 11. Identifying the installed version of the DDN Silicon Storage Appliance Manager software	56
Procedure 12. Installing the Silicon Storage Appliance Manager software for DDN devices	57
Procedure 13. Configuring the syslog on DDN storage devices	58
Procedure 14. Installing the SANtricity software	59
Procedure 15. Creating the boot RAID volume group for NetApp, Inc. Engenio devices	61
Procedure 16. Creating and configuring volumes for NetApp, Inc. Engenio devices	62
Procedure 17. Configuring zoning for a QLogic SANbox switch using QuickTools utility	64
Procedure 18. Creating a backup of your QLogic switch configuration	65
Procedure 19. Rebooting the SMW and verifying LUNs are recognized	66
Procedure 20. Shutting down the Cray system	68
Procedure 21. Updating the SMW software and configuration	68

	<i>Page</i>
Procedure 22. For Cray Cascade systems only: Updating the BC and CC firmware	71
Procedure 23. Updating the L0 and L1 firmware	74
Procedure 24. Confirming the SMW is communicating with system hardware	75
Procedure 25. Verifying hardware Programmable Interface Controllers (PICs)	76
Procedure 26. Updating SMW software on the boot root and shared root	77
Procedure 27. Changing the time zone for the SMW and the blade and cabinet controllers	81
Procedure 28. Starting the VNC server	85
Procedure 29. For workstation or laptop running Linux: Connecting to the VNC server through an <code>ssh</code> tunnel, using the <code>vncviewer -via</code> option	86
Procedure 30. For workstation or laptop running Linux: Connecting to the VNC server through an <code>ssh</code> tunnel	87
Procedure 31. For workstation or laptop running Mac OS X: Connecting to the VNC server through an <code>ssh</code> tunnel	87
Procedure 32. For workstation or laptop running Windows: Connecting to the VNC server through an <code>ssh</code> tunnel	88
Procedure 33. Setting up the SUSE firewall and IP tables	89
Procedure 34. Rack-mount SMW: Creating an SMW bootable backup drive	91
Procedure 35. Rack-mount SMW: Setting up the bootable backup drive as an alternate boot device	103
Procedure 36. Desk-side SMW: Creating an SMW bootable backup drive	107
Procedure 37. Desk-side SMW: Setting up the bootable backup drive as an alternate boot device	117
Procedure 38. Enabling an Integrated Dell Remote Access Controller (iDRAC6) on a rack-mount SMW	121
Procedure 39. Using the iDRAC6	122
Procedure 40. Rack-mount SMW: Installing the Dell Systems Management Tools and Documentation DVD	125
Procedure 41. Rack-mount SMW: Configuring the SMW to restore AC power after a power failure	126
Procedure 42. Rack-mount SMW: Replacing a failed LOGDISK or DBDISK disk drive	127

Examples

Example 1. LUN configuration dialog for DDN devices	52
Example 2. Final boot RAID LUN configuration for DDN devices	52

Tables

Table 1. Default IP Addresses	14
Table 2. Site-dependent Configuration Values	15
Table 3. Default System Passwords	16
Table 4. Identifying Linux <code>/dev</code> Names to Physical Drive Slots	21
Table 5. Identifying the <code>phy</code> in the Device Name Identifies the Drive Slot	21

Figures

Figure 1. Rack-mount SMW Rear Connections	17
Figure 2. Front View of the Rack-mount SMW Chassis Hard Drives	20

Page

Figure 3. Desk-side SMW: SMW Rear Connections 25

Introduction [1]

This guide contains procedures for the Cray System Management Workstation (SMW). Your SMW will be either the second-generation SMW, which is a rack-mount SMW, or a desk-side SMW. The following procedures are provided:

- Performing an initial installation of the SMW 7.0 release package; the SMW 7.0 software consists of the SMW base operating system, which is SUSE Linux Enterprise Server version 11 Service Pack 2 (SLES 11 SP2), and the Cray SMW software. See [Chapter 2, Installing the SMW Release Package on page 13](#).
- Configuring the boot RAID; see [Chapter 3, Configuring Boot RAID on page 49](#).
- Updating your Cray SMW software; see [Chapter 4, Updating or Upgrading Your Cray SMW Software on page 67](#). You must be running the SUSE Linux Enterprise Server version 11 Service Pack 2 (SLES 11 SP2) SMW base operating system and a release of SMW 6.0 or SMW 7.0 on your SMW in order to perform the procedures in this chapter.

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

A PDF version of this guide is available on the CrayDoc CD and through the CrayPort website at <http://crayport.cray.com>.

1.1 Related Publications

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

- *SMW Release Errata* and the *SMW README*, which are provided with your SMW release package
- *Cray Linux Environment (CLE) Software Release Overview* (S-2425), which is provided with your CLE release package
- *Cray Linux Environment (CLE) Software Release Overview Supplement* (S-2497), which is provided with your CLE release package as of its UP01 release package
- *Installing and Configuring Cray Linux Environment (CLE) Software* (S-2444), which is provided with your CLE release package
- *Managing System Software* (S-2393), which is provided with your CLE release package

- *Managing Lustre for the Cray Linux Environment (CLE)* (S-0010), which is provided with your Cray Linux Environment (CLE) operating system release package
- *Using and Configuring System Environment Data Collections (SEDC)* (S-2491), which is provided with your SMW release package
- *Cray Application Developer's Environment Installation Guide* (S-2465), which is provided with your Cray Application Developer's Environment (CADE) release package

1.2 Distribution Media

The Cray SMW 7.0 release distribution media includes the following DVDs:

- If you are performing an initial installation of the SMW software, the first DVD is labeled `Cray-SMWbase11SP2-yyyymmdd` and contains the SMW base operating system (SLES 11 SP2). The second is labeled `Cray SMW 7.0.UPxx` and contains the Cray SMW 7.0 software packages. All software is installed by means of scripts and RPM Package Manager (RPM) files.
- If you are updating your SMW software, use only the DVD labeled `Cray SMW 7.0.UPxx`, which contains the Cray SMW 7.0 software packages. All software is installed by means of scripts and RPM Package Manager (RPM) files. You must be running the SUSE Linux Enterprise Server version 11 Service Pack 2 (SLES 11 SP2) SMW base operating system and a release of SMW 6.0 or SMW 7.0 on your SMW in order to update to the `Cray SMW 7.0.UPxx` release.

Installing the SMW Release Package [2]

Follow these procedures to perform an initial software installation of the Cray System Management Workstation (SMW) 7.0.UPxx release package.

Note: Only the procedure to install the SMW base operating system differs for rack-mount SMWs and desk-side SMWs. All other sections and procedures in this chapter apply to all SMWs.

After completing this chapter, if you need to configure, format, or zone the boot RAID (redundant array of independent disks) system, see [Chapter 3, Configuring Boot RAID on page 49](#).

Note: Some examples are left-justified to fit better on the page. Left-justification has no special significance.

2.1 Preparing to Install the SMW Release Software

2.1.1 Before You Start the SMW Release Installation

- Read the *SMW Release Errata* and the *SMW README* provided with your SMW release package for any additional installation-related requirements, corrections to this installation guide, and other relevant information about the release package.
- Read the Field Notices (FNs) related to kernel security fixes to identify if there are any changes to this release package.
- Read this section and ensure that the prerequisites are satisfied before beginning an initial SMW software installation; verify that the network connections are in place (see [Network Connections on page 14](#)) and your configuration values are known (see [Configuration Values on page 14](#)).

Important: When you install the SMW release running under SLES 11 SP2, the SMW has a firewall enabled. `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 [Appendix C, Remote Access to the SMW on page 85](#).

2.1.1.1 Network Connections

The following network connections are required.

- The SMW must have the following private network connections:

eth0 – To the customer network

eth1 – To the Hardware Supervisory System (HSS) network

eth2 – Reserved

eth3 – To the boot node

Note: Ethernet port assignments are valid only after the SMW software installation completes (see [Rack-mount SMW: Installing the SMW Base Operating System on page 17](#)).

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

Important: Ensure you have protective covers for the Fibre Channel fibre optic cables and connectors. Before you perform the initial installation of the SLES software on your SMW, you will need to disconnect the SMW connection to the boot RAID.

2.1.1.2 Configuration Values

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

Table 1. Default IP Addresses

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

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

Table 2. Site-dependent Configuration Values

Description	Example	Actual Value
SMW hostname	xtsmw	
Domain	cray.com	
Aliases	cray-smw smw01	
Customer network IP address	192.168.78.68	
(Optional) For rack-mount SMW only: iDRAC6 IP address	cray-drac	
Customer network netmask	255.255.255.0	
Default gateway	192.168.78.1	
Domain names to search	us.cray.com mw.cray.com	
Nameserver IP address	10.0.73.30 10.0.17.16	
Timezone	US/Central	
NTP servers	ntpghost1 ntpghost2	
System Class	0, 1, 2, 3	
X dimension	1-64	
Y dimension	Cray XE and Cray XK systems: 1-16; Cray Cascade systems: 1-32	
Topology	0, 1, 2, 3, 4	

2.1.1.3 Passwords

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

Table 3. Default System Passwords

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

2.1.2 Shutting Down the Cray System



Warning: If your Cray system is currently installed, you **must** shut down your Cray system before proceeding with your SMW software installation.

Procedure 1. Shutting down the Cray system

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

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

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

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

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

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

Alternatively, you can execute these commands as `root` from the boot node to shutdown your system.

```
boot:~# xtshutdown -y
boot:~# shutdown -h now;exit
```

If you have a desk-side SMW, skip to [Desk-side SMW: Installing the SMW Base Operating System on page 25](#).

2.2 Rack-mount SMW: Installing the SMW Base Operating System

A full initial installation begins with installing the base operating system from the DVD on to the SMW ([Procedure 2 on page 17](#)). The base operating system is labeled `Cray-SMWbase11SP2-yyyyymmdd` on the DVD.


Procedure 2. Rack-mount SMW: Installing the base operating system

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

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

2. Disconnect the SMW connection to the boot RAID; disconnect the data cables and place protective covers on the fibre optic cables and connectors (if present).

Figure 1. Rack-mount 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** set-up screen appears.

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

A window listing the following appears:

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

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

A window listing the following appears.

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

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

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

8. Press the Esc key to exit the **Boot Sequence** window.

9. Press the Esc key again to exit the **Boot Settings** window.

10. Insert the base operating system DVD labeled
Cray-SMWbase11SP2-yyyymmdd into the CD/DVD drive.

11. Press the Esc key a final time to save changes and exit the BIOS set up. A screen listing exit options appears.

```
Save changes and exit
Discard changes and exit
Return to Setup
```

12. Ensure that **Save changes and exit** is highlighted. Then press the Enter key. The SMW resets automatically.

13. When the BIOS POST process begins again, **within 5 seconds, press the F11 key** after the following messages appear in the upper-right of the screen.

```

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

```

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

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

```

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

```

```

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

```

```

iDRAC6 Firmware Revision version: 1.54.15
Primary Backplane Firmware Revision 1.07

```

```

-----
                                IPv6 Settings
                                -----

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

                                IPv4 Settings
                                -----

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

```

The **iDRAC6 Configuration Utility** window appears.

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

The **Virtual Media Configuration** window appears.

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

```

Virtual Media ..... Detached

```

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

The **BIOS Boot Manager** screen appears.

19. Use the up-arrow and down-arrow keys to highlight the SATA Optical Drive entry.
20. Press the Enter key to boot from the installation DVD.
21. **Within 10 to 15 seconds** after the **SUSE Linux Enterprise Server** boot menu displays, use the down-arrow key to scroll down and select the **Cray SMW Initial Install Rackmount** option.
 - Boot from Hard Disk
 - Cray SMW Initial Install
 - **Cray SMW Initial Install Rackmount**
 - Cray SMW Upgrade Install
 - Repair Installed System
 - Rescue System
 - Check Installation Media

Then press the Enter key.

Otherwise, the first selection on this boot menu list is activated (that is, boot from the SMW's disk).

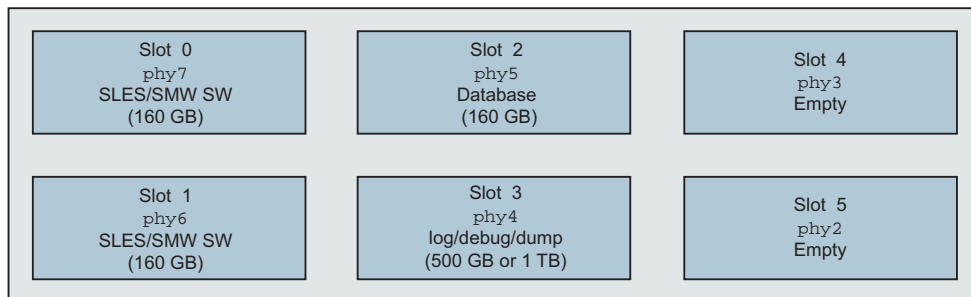
As the base installation progresses, the following phases appear on your screen:

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

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

22. Partition the drive for installation of the base OS and SMW software.
 - a. Map the physical drive slots to the Linux device names.

Figure 2. Front View of the Rack-mount SMW Chassis Hard Drives



- Slot 0: Target drive for installation of the base OS and SMW software.
- Slot 1: Will contain the backup installation of the base OS 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.

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

Table 4. Identifying Linux /dev Names to Physical Drive Slots

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

The identification of phy in the device name identifies the drive slot with the following mapping:

Table 5. Identifying the phy in the Device Name Identifies the Drive Slot

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

- Remove and recreate the phy7 partitions.

For each device that is phy7 **Device Path** (in our example, /dev/sdc1):

- Double-click on the **Device** name of that specific **Hard Disk** table entry, which opens the **Partition: /dev/sd...** table for that drive and which has **Edit**, **Resize** and **Delete** buttons below the table.
- Click on the **Delete** button. A pop-up window indicating **Really delete /dev/sd...** appears.

- 3) Verify the device, and click **Yes** to delete the existing disk partition on that drive. The specific **/dev/sd...** entry will be removed from the displayed list.

Note: Perform the above steps for all entries associated with the `phy7` device path until the displayed list is empty.

- c. Create new partitions for `swap` and `root` for the `phy7` device sized for your system. This is the target boot device for the OS installation.

- 1) Click on the **Add** button at the bottom of the **Hard Disk: /dev/sd...** window, which opens the **Add Partition on /dev/sd...** screen.
- 2) Select **Primary Partition**, then click **Next**.

The **New Partition Size** screen displays.

- 3) Select **Custom size**, then enter the swap partition size into the box.

This swap partition size should be the same as the total system RAM memory on the SMW.

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

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

Then click **Next**.

The **Formatting Options** and **Mounting Options** screen displays.

- 4) Select **Format partition** and select **Swap** from the pull-down menu under **File system**.
- 5) Select **Mount partition** and select **swap** from the pull-down menu under **Mount Point**.
- 6) Click on the **Finish** button.

The **Hard Disk: /dev/sd...** window displays again, but with the new partition entry.

- 7) Click on the **Add** button at the bottom of the window, which opens the **Add Partition on /dev/sd...** screen.
- 8) Select **Primary Partition**, then click **Next**.

The **New Partition Size** screen displays.

- 9) Select **Maximum Size**, then click **Next**.

The **Formatting Options** and **Mounting Options** screen displays.

- 10) Select **Format partition** and select **Ext3** from the pull-down menu under **File system**.
 - 11) Select **Mount partition** and select **/** from the pull-down menu under **Mount Point**.
 - 12) Click **Finish**.
- d. Partitioning of the target boot device for the OS installation has been completed. On the **Expert Partitioner** screen, click **Accept** to accept the changes.

A pop-up window stating the following will be displayed:

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

- e. Click **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.

- f. Confirm the disk partitions.

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

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

If further changes are necessary, click on the **Partitioning** section header and return to the **Expert Partitioner** screen to make adjustments.

23. Confirm the language for the SMW.

English (US) is the default language. To change the primary language, click on the **Language** heading in the **Installation Settings** screen. The **Languages** window opens. Select your language choice from the drop-down menu. You can also select multiple secondary languages. Then click on **Accept** at the bottom of the window.

24. Continue the Installation.

On the **Installation Settings** screen, click on **Install** at the bottom of the screen. The **Confirm Installation** window appears. If you want to check or change settings, click on **Back**; otherwise, click on **Install** to confirm and to install the operating system.

Note: 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. Next, the **Network Settings** window for configuring the customer network appears.

- a. Select the entry labeled **eth0 Customer Network Ethernet**.
- b. On the **Overview** tab, click on **edit** and do the following:
 - 1) Enter the IP address.
 - 2) Enter the subnet mask.
 - 3) Enter the short host name and click **Next**.
- c. Click on the **Hostname/DNS** tab. In the **Hostname/DNS and Name Server Configuration** window, enter the host name, the domain name, the name server values, and the domain names to search. You must enter the host name and domain name separately; for example:

Hostname	Domain Name
smwhost	my.domain.com

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

- d. Click on the **Routing** tab. In the **Routing** window, enter the Default Gateway IP address.

Then click on **OK**.

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

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

Note: At this point, the SMW comes up the rest of the way to multiuser mode for login on the console.

The SMW base operating system, `Cray-SMWbase11SP2`, is now installed. Continue with the next step.

27. Remove the protective covers from the fiber optic cables and connectors, clean the ends of the cables and connectors, and reconnect the data cables.

28. Re-seat the previously ejected disk drives.

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

```
smw:~ # reboot
```


30. (Optional) To enable an Integrated Dell Remote Access Controller 6 (iDRAC6), see [Appendix I, Enabling an Integrated Dell Remote Access Controller \(iDRAC6\) on a Rack-mount SMW on page 121](#).

Skip to [Installing and Configuring the SMW Software Packages on page 30](#) to install and configure the SMW software packages.

2.3 Desk-side SMW: Installing the SMW Base Operating System

A full initial installation begins with installing the base operating system from the DVD on to the SMW ([Procedure 3 on page 25](#)). The base operating system is labeled `Cray-SMWbase11SP2-yyyymmdd` on the DVD.


Procedure 3. Desk-side SMW: Installing the base operating system

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

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

2. Disconnect the SMW connection to the boot RAID; disconnect the data cables and place protective covers on the fibre optic cables and connectors (if present).

Figure 3. Desk-side SMW: SMW Rear Connections



3. Power up the SMW. When the BIOS power-on self-test (POST) process begins, **quickly press the Del 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 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 ordered as follows:

```
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 [Procedure 4](#) (see [Installing and Configuring the SMW Software Packages on page 30](#)). If this happens, the only solution is to return to [Procedure 3 on page 25, step 3](#), and repeat every step in both [Procedure 3](#) and [Procedure 4](#). Therefore, ensure that the drive order is set correctly at this time.

9. Press the Esc key to return to the **Boot** menu. Then press the left arrow to select the **Main** menu. On the **Main** menu, the **System Time** option should already be selected. Set the system date and time to the time at the Coordinated Universal Time (UTC) time zone. This time can be found by executing the command `date -u` on any system.

10. Insert the base operating system DVD labeled Cray-SMWbase11SP2-yyyymmdd into the CD/DVD drive.
11. Press the F10 key to save the BIOS changes. A dialog box displays the question **Save configuration changes and exit setup?**.
 - Press the Enter key to save changes, exit the setup utility, and automatically reboot the SMW.
 - Alternatively, press Tab to highlight Cancel and then press the Enter key to cancel acceptance of BIOS changes, so that you can 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-backgrounded 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.

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

Then press the Enter key.

Otherwise, the first selection on this boot menu list is activated (that is, boot from the SMW's disk).

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

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

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

15. Confirm which disk partitions are marked for use as swap on the SMW. 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

There are lines similar to the following line that show the primary SMW disk partition to be used for swap:

- Create swap partition /dev/sda1

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

- a. To designate that a disk partition not be used for swap on the SMW, click on the **Partitioning** heading, which opens the **Expert Partitioner** window.
 - b. Click on the + sign next to **Hard Disks** to display the hard disks that are connected.
 - c. Click on the + sign next to a disk device that incorrectly has a swap partition in the **Mount Point** column.
 - d. Click on **sdb1** (disk device) to open the **Edit Partition** window.
 - e. Click on **Edit** at the bottom of the window.
 - f. Click on **Do not mount partition**.
 - g. Click on **Finish** in the **Edit Partition** window.
 - h. Click on **Accept**.
16. Confirm the language for the SMW.

English (US) is the default language. To change the primary language, click on the **Language** heading in the **Installation Settings** window. The **Languages** window opens. Select your language choice from the drop-down menu. You can also select multiple secondary languages. Then click on **Accept** at the bottom of the window.

17. Continue the Installation.

On the **Installation Settings** window, click on **Install** at the bottom of the window. The **Confirm Installation** window appears. If you want to check or change settings, click on **Back**; otherwise, click on **Install** to confirm and to install the operating system.

Note: The installation runs for approximately 30 minutes. The process automatically reboots the SMW from the hard disk, and the installation continues with system configuration.

18. Next, the **Network Settings** window for configuring the network appears.

- a. Select the entry labeled **eth0 Customer Network Ethernet**.
- b. On the **Overview** tab, click on **edit** and do the following:
 - 1) Enter the IP address.
 - 2) Enter the subnet mask.
 - 3) Enter the short host name and click **next**.
- c. Click on the **Hostname/DNS** tab. In the **Hostname/DNS and Name Server Configuration** window, enter the host name, the domain name, the name server values, and the domain names to search. You must enter the host name and domain name separately; for example:

Hostname	Domain Name
smwhost	my.domain.com

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

- d. Click on the **Routing** tab. In the **Routing** window, enter the Default Gateway IP address.

Then click on **OK**.

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

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

Note: 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. Continue with the next step.

20. Remove the protective covers from the fiber optic cables and connectors, clean the ends of the cables and connectors, and reconnect the data cables.

2.4 Installing and Configuring the SMW Software Packages



Warning: You **must** shut down your Cray system before you install the Cray SMW software packages. To shut down your system, see [Procedure 1 on page 16](#).

Procedure 4. Installing and configuring SMW software

Note: This procedure takes approximately 30 minutes.

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

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

2. If the base operating system DVD (`Cray-SMWbase11SP2-yyyyymmdd`) is still in the CD/DVD drive, eject it.

```
smw:~ # eject
```

3. Place the Cray SMW 7.0.UPxx Software DVD in the drive and mount it.

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

Note: If you experience problems mounting the DVD after the initial Linux installation, reboot the SMW and begin this procedure again.

4. Copy the `SMWinstall.conf` file from `/media/cdrom` to `/home/crayadm` and modify it for your site.

The `SMWinstall.conf` file contains settings for the system interconnection network type; the name of the local Network Time Protocol (NTP) servers; configuration settings for SMWs that have additional disks and for an SMW of a Cray XE6m, Cray XK6m, or Cray XK7m mid-range system; configuration settings for the Cray Lightweight Log Manager (LLM); and a setting to enable or disable the Linux `sar` service on the SMW.

Important: For all systems:

- In the `SMWinstall.conf` file, you must specify the system interconnection network type for your system; there is no default setting.
- The `xttrim` utility is used to manage logs on the SMW. As of the SMW 7.0 release, the `xttrim` utility replaced the `xtclean_logs` utility, so the `ENABLE_XTCLEAN_LOGS` variable in the `SMWinstall.conf` file must be set to `no`.

After your SMW software installation is complete, review the `xttrim.conf` configuration file to ensure the settings are consistent with your site's log maintenance policies.

- By default, the `SMWinstall` program enables LLM. The `SMWinstall` program automatically updates the `llm.conf` configuration file during the SMW software installation or upgrade. If you make changes to the `llm.conf` file provided with the release package, you must be reapply your changes after an SMW software installation or upgrade. Your `llm.conf` configuration file will be saved as `/etc/opt/cray/llm/llm.conf.rpm` as part of the update process.
- By default, the Linux `sar` service on the SMW is enabled.
- Controller log forwarding is enabled by default.

Important: For rack-mount SMWs: The `LOGDISK` and `DBDISK` variables **must** be set.

- Use the `LOGDISK` variable to specify the disk device name to be used for logging (`/var/opt/cray/log`, `/var/opt/cray/dump`, and `/var/opt/cray/debug`). The entire disk will be formatted for this use. On a rack-mount SMW, this will be the `phy4` disk.
- Use the `DBDISK` variable to specify the disk device name to be used for the SMW HSS database (`/var/lib/mysql`). The entire disk will be formatted for this use. On a rack-mount SMW, this will be the `phy5` disk.
- `LOGDISK` must be larger than `DBDISK`.
- The `LOGDISK` and `DBDISK` variables must be persistent disk device names (see [Figure 2](#), [Table 4](#), and [Table 5](#)).

To use the default persistent disk device names, uncomment the following lines in the file:

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

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

If you have a nonstandard configuration, you will need to modify the persistent disk device names for the LOGDISK and DBDISK variables.

```
smw:~ # cp /media/cdrom/SMWinstall.conf /home/crayadm
smw:~ # chmod 644 /home/crayadm/SMWinstall.conf
smw:~ # vi /home/crayadm/SMWinstall.conf
```

Note: To reconfigure a rack-mount SMW's LOGDISK or DBDISK when you need to replace a failed drive with new drive, see [Appendix K, Rack-mount SMW: Replacing a Failed LOGDISK or DBDISK Disk Drive on page 127](#).

Important: For Cray XE6m, Cray XK6m, or Cray XK7m systems: The MCLASS and NumChassis variables **must** be set.

Note: These configuration settings do **not** apply to Cray XE6m-200, Cray XK6m-200, or Cray XK7m-200 systems.

MCLASS For a Cray XE6m, Cray XK6m, or Cray XK7m system only: MCLASS must be set to TRUE. For all other Cray systems, including Cray XE6m-200, Cray XK6m-200, and Cray XK7m-200 systems, MCLASS must be set to FALSE.

Default setting is FALSE.

NumChassis For a Cray XE6m, Cray XK6m, or Cray XK7m system only: Uncomment the NumChassis line; NumChassis 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.

5. 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 displays to the console. The SMWinstall installation script also creates log files in /var/adm/cray/logs.

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

6. Reboot the SMW.

```
smw:~ # reboot
```

7. (Optional) If you want to set up the Virtual Network Computing (VNC) server, see [Appendix C, Remote Access to the SMW on page 85](#).



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

Note: Henceforth, whenever the SMW presents a session selection menu, always select SUSE Linux Enterprise Server 11 SP2.

2.5 Discovering the Cray System Hardware and Powering Up the Full System

Procedure 5. Bootstrapping hardware discovery

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

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

2. The `xtdiscover` command you execute in this step needs to collect some basic information in order to bootstrap the hardware discovery process.

When started, `xtdiscover` warns that changes will be made and asks for a confirmation of aborting or continuing. You will need to enter information when you see:

```
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]
Boot node name [c0-0c0s0n1]
SDB node name [c0-0c0s2n1] (for Cray XE or Cray XK systems)
SDB node name [c0-0c0s1n1] (For Cray Cascade systems)
Pathname to the default boot image [/raw0]
```

Note:

- The default boot node provided by `xtdiscover` is `c0-0c0s0n1` for a Gemini network (Cray XE or a Cray XK systems) and for an Aries network (Cray Cascade systems).
- The default SDB node provided by `xtdiscover` is `c0-0c0s2n1` for a Gemini network (Cray XE or a Cray XK systems) and is `c0-0c0s1n1` for an Aries network (Cray Cascade systems).

By default, the `xtdiscover` command uses the `/opt/cray/hss/default/etc/xtdiscover.ini` file. (When you perform an SMW update, the process will not overwrite the `xtdiscover.ini` file; the new version of the file will be placed in `xtdiscover.ini.dist`.)

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

Enter the following command.

Note: If the `xtdiscover` command fails with the message, The following cabinets were not detected by heartbeat:, then you need to physically power cycle the cabinet controller, and then retry 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
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.

Executing command '/opt/cray/hss/default/bin/fm -w -t l1 s0'. Output follows...

c0-0: ok

Command succeeded.
```

Executing command '/opt/cray/hss/default/bin/xtrsh -m '-[0-9]{1,}\$' -l root -s reboot'. \

Output follows...

c0-0 : :::0
c0-0 : OK

Command succeeded.

Rebooting Lls.
Waiting for Lls to finish rebooting.
....
All Lls are rebooted and ready

Suspending State Manager to continue with cabinet discovery...
Suspend successful.

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

Clearing database...done.
Storing base cabinet data...done.
Resuming State Manager for power-up...Resume successful.

Discovery Phase 1 of 3 complete.

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

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

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

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

Powering on cabinet...
1 cabinet will be powered on:
[1 out of 1]
Cabinets powered on.
Storing base cabinet data...done.

Skipping bounce and discover phases 2 & 3.

Restarting RSMS daemons for normal operation:
Stopping RSMS services: cm sedc_manager bm sm erd done
Starting RSMS services: erd sm nm bm sedc_manager cm
Flushing and installing cabinet routes...done.
Checking state manager availability...done.
Partition p0 is active but not booted; skipping background commands. done
Done.

Discovery complete
***** xtdiscover finished *****

3. For a Cray Cascade system (Aries network) SMW only; otherwise, skip to [step 4](#).

a. Update the controller boot image.

```
smw:~ # cat /opt/cray/hss/default/etc/crms-release
HSS-CRMS Tue Oct 23 01:31:09 CDT 2012 on hssbld0 by bwdev
lsb-cray-hss-7.0.0-1.0700.29435.242
smw:~ # hss_make_default_initrd /opt/cray/hss-images/master/7.0.0-1.0700.29435.242
::: Verifying base RPM list to the manifest
::: Installing filesystem hierarchy
::: Installing busybox
::: Installing base RPMs
::: Installing ssh
::: Removing /etc/securetty installed by the pam RPM
::: Installing sshfs
::: Installing rsh
::: Modifying /etc/pam.d/rlogin to remove securetty checking
::: Modifying /etc/pam.d/rsh to remove rhosts and nologin checking (Bug #779466)
::: Installing rsync
::: Installing lsof
::: Installing atftp
::: Installing tcpdump
::: Installing ethtool
::: Installing syslog-ng
::: Installing logrotate
::: Installing ntp/ntpd
::: Installing strace
::: Installing screen
::: Installing minicom
::: Installing ppp
::: Installing mtd-utils
::: Installing /init
::: Installing file.rpm
::: Installing libgmodule
::: Installing Midnight Commander
::: Installing cray-viper
::: Installing spread
::: Installing coreboot-utils
::: Clearing init.d to be replaced by cray-hss32-filesystem
::: Creating initial etc files needed for root creation
::: Installing Cray kernel
::: Installing latest Cray kernel modules
::: Clearing select /boot items
::: Installing boot-parameters
::: Installing cray-hss32-scripts
::: Installing lsb-cray-hss-controllers
::: Installing cray-libconfig
::: Installing cray-bdm
::: Installing cray-play_xsvf
::: Removing unwanted files from the root

=====
The new initrd hierarchy is now in /opt/cray/hss-images/master/7.0.0-1.0700.29435.242.

Running hssclone.
Image Clone Complete: /opt/cray/hss-images/default
Running hsspackage.
```

```

copying image
copying modules
running depmod
creating load file: /opt/cray/hss-images/default/HSS32.load
compressing initrd.img
Creating pxelinux.0 symlink
Running hssbootlink.
linking /opt/cray/hss-images/default/HSS32/bzImage-3.0.13-0.27.1_1.0700.6694-cray_hss32 \
/opt/tftpboot/bzImage
linking /opt/cray/hss-images/default/HSS32/parameters \
/opt/tftpboot/pxelinux.cfg/default
linking /opt/cray/hss-images/default/HSS32/initrd.img /opt/tftpboot/initrd.img
smw:~ # exit

```

- b. Power down the system.

```
crayadm@smw:~> xtcli power down s0
```

- c. For each cabinet controller, execute the following command, and then ensure the cabinet controller is up:

```

crayadm@smw:~> xtlogin c0-0
(CC) c0-0 / # /sbin/reboot
crayadm@smw:~> xtalive -l cc c0-0

```

- d. Power up the system.

```
crayadm@smw:~> xtcli power up s0
```

4. **For a Cray XE or a Cray XK system (Gemini network) SMW only; otherwise, skip to [Procedure 6 on page 39](#).**

- 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]$" -f /tmp/xtrsh.Sk5yg9Uz -l root -s "if [ -x /etc/loadnor ]; \
then /etc/loadnor; fi ; if md5sum /dev/mtd/4 | grep fec74d6d797dd8b36e68282ad43d0773; \
then echo OLDBIOS; fi"
There is one L1 not up-to-date.
spawn fm -qRM -t l0 s0
It took 2 seconds for 'query' to complete to all L0s.
xtrsh -m s -f /tmp/xtrsh.6ef3luGH -l root -s "if [ -x /etc/loadnor ]; \
then /etc/loadnor; fi ; if md5sum /dev/mtd/4 | grep fec74d6d797dd8b36e68282ad43d0773; \
then echo OLDBIOS; fi"
There are 24 L0s not up-to-date.
#####
# Attempt #1 #
#####
spawn fm -RM -t l1 c0-0
It took 157 seconds (2 minutes, 37 seconds) for 'flash' to complete to all L1s.
There is one L1 not up-to-date.
xtrsh -m "-[0-7]$" -f /tmp/xtrsh.b0zXZERW -l root -s "reboot"
Sleeping for 180 seconds...
It took 0 seconds for 'ping' to complete to all L1s.
It took 0 seconds for L1s to become available.
spawn fm -RM -t l0 c0-0c0s0 c0-0c0s1 c0-0c0s2 c0-0c0s3 c0-0c0s4 c0-0c0s5 c0-0c0s6 c0-0c0s7 \
c0-0c1s0 c0-0c1s1 c0-0c1s2 c0-0c1s3 c0-0c1s4 c0-0c1s5 c0-0c1s6 c0-0c1s7 c0-0c2s0 c0-0c2s1 \
c0-0c2s2 c0-0c2s3 c0-0c2s4 c0-0c2s5 c0-0c2s6 c0-0c2s7
It took 180 seconds (3 minutes, 0 seconds) for 'flash' to complete to all L0s.
There are 24 L0s not up-to-date.
.
.
.
Everyone is properly flashed.
It took 871 seconds (14 minutes, 31 seconds) to run this application.
crayadm@smw:~>
```

- c. If all the L0s are running a compatible BIOS firmware, then `xtflash` completes successfully, and you can continue to [Procedure 6 on page 39](#). However, If the L0s do not get flashed with the updated BIOS firmware, `xtflash` will halt and display a list of those L0s that need to be updated. These L0s must have their BIOS reflashed before `xtflash` will continue. Only perform the BIOS flash on the components that require it.

If `xtflash` halts and displays a list of those L0s that need to be updated, Cray recommends reflashing the Cray L0s in groups of a few cabinets at a time, rather than the entire system at once. If there is a catastrophic failure during the BIOS flashing procedure, the only way to recover the BIOS is to reflash each BIOS chip by hand. By flashing only a few cabinets at a time,

you limit the risk that a catastrophic failure can cause. In addition, if you flash the BIOS on an L1/L0 that already has a compatible BIOS (like an SIO), then you must power cycle the L1/L0 after flashing it.

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

- 1) For example:

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

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

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

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

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

- d. Continue with [Procedure 6 on page 39](#).

Procedure 6. Discovering the Cray system hardware

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

Remember that the `xtdiscover` command needs to collect some basic information in order to bootstrap the hardware discovery process, warns that changes will be made, and asks for a confirmation of aborting or continuing. You need the X and Y cabinet sizes as well as topology from [Table 2 in Configuration Values on page 14](#) to accurately discover your hardware. The `xtdiscover` command will also prompt you to execute the `xtbounce` command in a separate window and then continue.

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

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

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

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

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

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

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

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

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

Enter your system's network topology class [0]: 0
Setting topology class to 0

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

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

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

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

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

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

Discovery Phase 1 of 3 complete.

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

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

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


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

Powering on cabinet...

1 cabinet will be powered on:

[1 out of 1]

Cabinets powered on.

Discovering component phase 2 (blade) state:

[24 out of 24]

Finished discovering component phase 2 (blade) state.

Discovering component phase 2 (blade) attributes:

[24 out of 24]

Finished discovering component phase 2 (blade) attributes.

Verifying phase 2 configuration...done.

Summary of blades discovered:

Total:	24	Service:	11	Empty:	0	Disabled:	0
--------	----	----------	----	--------	---	-----------	---

Storing attribute data...done.

Discovery Phase 2 of 3 complete.

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

Resuming State Manager for bounce...Resume successful.

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

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

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

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

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

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

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

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

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

Discovering component phase 3 (blade/node) attributes:
[24 out of 24]
Finished discovering component phase 3 attributes.
Verifying configuration...
INFO: 24 newly discovered components were added.
INFO: 24 components in previous configuration were deleted.
INFO: Added the following hardware:
 24 slots
INFO: Removed the following hardware:
 24 slots
INFO: Configuration change details are in
/opt/cray/hss/default/etc/xtdiscover-config-changes.diff

done.
Storing component attribute data...done.
Updating component location history...done.

Restarting RSMS daemons for normal operation:
Stopping RSMS services: cm sedc_manager bm nm sm erd done
Starting RSMS services: erd sm nm bm sedc_manager cm
Flushing and installing cabinet routes...done.
Checking state manager availability...done.
Partition p0 is active but not booted; skipping background commands.

done

Done.

Discovery complete
***** xtdiscover finished *****

4. **For a Cray Cascade system (Aries network) SMW only; otherwise, skip to [step 5](#).** Update the firmware.

- a. Exit from root.

```
smw:~ # exit
```

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



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

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

Note: If partitions are enabled and `p0` is not, then you must specify each active partition by using `pN`, where *N* designates the partition name:

```
crayadm@smw:~> xtzap -r -v p1
```

- c. 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)chia_fpga (CHIA FPGA)
```

This is a list of all blade level components:

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

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

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

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

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

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

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

- d. Execute the `xtzap -r -v s0` (or `pN`) command again; all firmware revisions should report correctly, except `node_bios`; `node_bios` will display as "NOT_FOUND" until after you complete [step 4.e](#) (the `xtbounce --linktune` command completes).

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

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

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

5. As user `root`, place the Cray SMW 7.0.UPxx Software DVD in the drive and mount it again.

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

6. Execute the `SMWconfig` command to configure MySQL, and start Cray Management Services (CMS). When you are prompted for your old root MySQL database password, press the Enter key (that is, an empty string), and then set your new root MySQL database password to whatever you want.

Note: If you do not enter your new root MySQL database password when prompted for it but press the Enter key (that is, an empty string) instead, you will need to enter a new root MySQL database password at the DB Root Password: prompt when `SMWconfig` spawns `mzinit`.

The SMWconfig command may 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
17:00:33 Finished checking command line options
17:00:33 Installation output will be captured in \
/var/adm/cray/logs/SMWconfig.stdout.7814
17:00:33 Installation errors (stderr) will be captured in \
/var/adm/cray/logs/SMWconfig.stderr.7814
17:00:33 Beginning second part of installation.
17:00:33 Detected hss dir: /opt/cray/hss/default/bin.
17:00:33 Detected network type: Gemini
.
17:00:33 Checking SMW version...
17:00:33 SMW version detected: v6.0
17:00:33 Detected hss dir: /opt/cray/hss/default/bin.
17:00:34 Initial run of SMWconfig.
17:00:34 SMW Configuration File: /home/crayadm/SMWinstall.conf
17:00:34 Reading configuration.
17:00:34 Opening Installation Configuration File: /home/crayadm/SMWinstall.conf
17:00:34 Finishing with read_config_file
```

```
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.
17:00:41 Initializing the CMS log manager.
17:00:41 Beginning to write the external mzinit expect script.
17:00:41 Finished writing the external mzinit expect script.
spawn mzinit
DB Root Password: (If prompted for one, enter the new root MySQL database password.)
17:00:44 The CMS log manager has been initialized.
17:00:44 Beginning clean up...
17:00:44 Finished cleaning up.
17:00:44 SMWconfig has completed.
```

7. If you configure postfix on your SMW, complete this step; otherwise, skip to [step 8](#).

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

```
MAIL_CREATE_CONFIG="no"
```

8. Unmount the Cray SMW 7.0.UPxx Software DVD.

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

2.6 Confirming the SMW is Communicating with the System Hardware

Procedure 7. Confirming the SMW is communicating with the system hardware

Ensure that the SMW is functional by executing the following three commands.

1. Execute the `xthwinv` command on the entire system (`s0`) and confirm that the processor type and counts that are reported match the configuration for your system. For additional information, see the `xthwinv(8)` man page.

```
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` command 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]
.				
.				
.				

```
crayadm@smw:~>
```

3. Execute the `xtmcinfo -t -u s0` command to retrieve microcontroller information from the cabinet control processors and the blade control processors.

```
crayadm@smw:~> xtmcinfo -t -u s0
c0-0c0s0 - Wed Apr 20 01:27:27 CDT 2012
c0-0c0s1 - Wed Apr 20 01:27:27 CDT 2012
c0-0c0s2 - Wed Apr 20 01:27:27 CDT 2012
c0-0c0s3 - Wed Apr 20 01:27:27 CDT 2012
c0-0c0s4 - Wed Apr 20 01:27:27 CDT 2012
c0-0c0s5 - Wed Apr 20 01:27:27 CDT 2012
c0-0c0s6 - Wed Apr 20 01:27:27 CDT 2012
c0-0c0s7 - Wed Apr 20 01:27:27 CDT 2012
.
.
.
c0-0c2s6 - Wed Apr 20 01:27:27 CDT 2012
c0-0c2s7 - Wed Apr 20 01:27:27 CDT 2012

How long have they been up...
c0-0c0s0 - 14:11:49 up 23 min, load average: 0.00, 0.01, 0.03
c0-0c0s1 - 14:11:49 up 23 min, load average: 0.53, 0.13, 0.06
c0-0c0s2 - 14:11:50 up 23 min, load average: 0.00, 0.00, 0.02
c0-0c0s3 - 14:11:50 up 23 min, load average: 0.14, 0.03, 0.01
c0-0c0s4 - 14:11:50 up 23 min, load average: 0.00, 0.01, 0.03
c0-0c0s5 - 14:11:50 up 23 min, load average: 0.01, 0.10, 0.12
c0-0c0s6 - 14:11:50 up 23 min, load average: 0.04, 0.10, 0.10
c0-0c0s7 - 14:11:50 up 23 min, load average: 0.01, 0.07, 0.08
.
.
.
c0-0c2s6 - 14:11:52 up 23 min, load average: 0.03, 0.11, 0.09
c0-0c2s7 - 14:11:53 up 23 min, load average: 0.01, 0.09, 0.08
crayadm@smw:~>
```

2.7 For Cray Systems with Cray XK Blades: Updating the Cray XK Accelerator Firmware

Important: For Cray systems with Cray XK blades: After completing the SMW software installation, you must contact your Cray service personnel to update the Cray XK accelerator firmware so that it is synchronized with the CLE release of the Cray XK driver. These steps **must** be done **before** installing CLE or booting the Cray system and can be run even if your site does not install the latest CLE release package.

2.8 Changing Default SMW Passwords After Completing Installation

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

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

2.9 Setting Up the SUSE Firewall and IP Tables

To configure your firewall, see [Appendix D, Setting Up the SUSE Firewall and IP Tables on page 89](#).

2.10 Collecting Additional Data During a System Dump Session

The `xtdumpsys` command collects information from a Cray system that is failing or has failed, has crashed, or is hung. Cray service personnel use this information to analyze the data to attempt to determine the cause of the problem.

As shipped, `xtdumpsys` does not collect any logs from various service nodes. Without these logs, Cray service personnel will not be able to analyze many types of problems. Cray strongly recommends that you copy `/opt/cray/hss/default/etc/example.xtdumpsys-plugin` to `~crayadm/.xtdumpsys-plugin` and customize it for your site, in order to have `xtdumpsys` collect this additional information automatically. You will need to modify it to match the way your site names its systems, allows logins from the SMW to the boot node and from the boot node to other service nodes, and stores system dumps long-term.

2.11 Backing Up Your Newly-installed SMW Software

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

- If you have a rack-mount SMW, see [Appendix E, Rack-mount SMW: Creating a Bootable Backup Drive on page 91](#).
- If you have a desk-side SMW, see [Appendix G, Desk-side SMW: Creating a Bootable Backup Drive on page 107](#).

Configuring Boot RAID [3]

3.1 Prerequisites and Assumptions for Configuring the Boot RAID

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

Cray provides support for system boot RAID from two different vendors: Data Direct Networks, Inc. (DDN) and NetApp, Inc. If you are configuring a DDN boot RAID, follow the procedures in [Configuring and Zoning the Boot RAID Using a DDN Storage System on page 50](#). If you are configuring a NetApp, Inc. Engenio boot RAID, follow the procedures in [Configuring the Boot RAID Using a NetApp, Inc. Engenio Storage System on page 59](#).

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.

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

This chapter presupposes that:

- 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) connection to the boot RAID or to a FC switch.
- The boot nodes have a FC connection to the boot RAID or to a FC switch.
- The SDB node has a FC connection to the boot RAID or to a FC switch.
- If a dedicated syslog node is configured, it has a FC connection to the boot RAID or to a FC switch.

- The DDN RAID uses LUNs (Logical Units); the NetApp, Inc. RAID uses volumes.

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

3.2 Configuring and Zoning the Boot RAID Using a DDN Storage System

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

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

3.2.1 Using the Command-line Interface (`telnet`) to Configure a DDN Storage System for Your Boot RAID

3.2.1.1 Configuring the Boot RAID Password When Using DDN Devices

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`. Cray recommends that you `telnet` to the RAID controller and change the passwords for `admin` and `user`.

Procedure 8. Configuring the boot RAID password when using DDN devices

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

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

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

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

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

```

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

Administrative user name 'admin' unchanged.

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

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

General user name 'user' unchanged.

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

Password for general user 'user' unchanged.

Committing changes.
CAB01-L1:

```

3.2.1.2 Configuring the LUNs for DDN Devices

Configure the boot RAID with at least six LUNs to support the various system management file systems.

Procedure 9. Configuring the LUNs for DDN devices

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

```

CAB01-L1: lun

Logical Unit Status

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

```

2. (Optional) Use the `lun` command to delete preexisting LUNs.

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

where *x* is the LUN to be deleted.

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

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

where *x* is the LUN to be configured. This command initiates a dialog similar to the following example.

Example 1. LUN configuration dialog for DDN devices

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

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

4. Repeat [step 3](#) for each LUN that you want to configure and format. The final LUN configuration resembles this example.

Example 2. Final boot RAID LUN configuration for DDN devices

```
CAB01-L1: lun
```

```
Logical Unit Status
```

LUN	Label	Owner	Status	Capacity (Mbytes)	Block Size	Tiers	Tier list
0	bootroot0	1	Ready	40000	4096	1	1
1	bootroot1	1	Ready	40000	4096	1	2
2	shroot0	1	Ready	280000	4096	1	3
3	sdb	1	Ready	80000	4096	1	2
4	syslog	1	Ready	80000	4096	1	1
5	shroot1	1	Ready	280000	4096	1	4
6	bootroot2	1	Ready	40000	4096	1	3
7	shroot2	1	Ready	280000	4096	1	1

3.2.1.3 Zoning the LUNs for DDN Devices

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

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

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

Procedure 10. Zoning the LUNs for DDN devices

1. Display the current zoning summary.

```
CAB01-L1: zoning
```

```
Port Zoning Summary:
```

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

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

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

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

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

This command returns the following prompt:

```
Enter the new LUN zoning for host port 1.
```

```
Enter the unique LUN mapping, as follows:
```

```
G.1  GROUP.LUN number
P    Place-holder
R    Before GROUP.LUN to indicate Read-Only
N    Clear current assignment
<cr> No change
E    Exit command
?    Display detailed help text
```

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

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

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

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

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

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

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

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

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

```
External LUN 0: is not mapped. Enter new internal LUN: n
External LUN 1: is not mapped. Enter new internal LUN: n
External LUN 2: is not mapped. Enter new internal LUN: n
External LUN 3: is not mapped. Enter new internal LUN: n
External LUN 4: is not mapped. Enter new internal LUN: n
External LUN 5: is not mapped. Enter new internal LUN: n
External LUN 6: is not mapped. Enter new internal LUN: 6
External LUN 7: is not mapped. Enter new internal LUN: 7
External LUN 8: is not mapped. Enter new internal LUN: 8
External LUN 9: is not mapped. Enter new internal LUN: e
*** Host Port 3: zoning has been updated! ***
                                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 you perform maintenance on disk subsystems. If you are running LUN verification and swap or move a back-end channel cable, you can disrupt an entire channel of drives.

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

```
CAB01-L1: zoning
                                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 you have finished zoning the LUNs, close the telnet connection to the RAID controller and return to the SMW.

```
CAB01-L1: logout
```

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

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

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

```
smw:~ # reboot
```

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

3.2.2 Installing the Silicon Storage Appliance Manager GUI to Configure a DDN Storage System for Your Boot RAID

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

3.2.2.1 Installing the DDN Silicon Storage Appliance Manager

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

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

Procedure 11. Identifying the installed version of the DDN Silicon Storage Appliance Manager software

1. Start the GUI.

```
# /usr/local/ddn/Silicon_Storage_Appliance_Manager
```

Or

```
# /opt/Silicon_Storage_Appliance_Manager/Silicon_Storage_Appliance_Manager
```

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

Procedure 12. Installing the Silicon Storage Appliance Manager software for DDN devices

1. Log on to the master node as `root` user. Forward your X Window display from the Cray system to your workstation.

```
# ssh -X
```

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

```
# /usr/local/ddn/DDN_version_Linux_VM.bin
```

The **Introduction** window displays.

4. Click **Next**.

The **License Agreement** window displays.

5. Accept the License Agreement and click **Next**.

The **Choose Uninstall Folder** window displays.

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

The **Choose Java Virtual Machine** window opens.

7. Choose a Java Virtual Machine and click **Next**.

The **Choose Install Folder** window opens.

8. Choose the **install folder** and click **Next**.

The **Choose Link Folder** window opens.

9. Choose the **link folder** and click **Next**.

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

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

When the installation completes, the **Install Complete** window opens.

11. The installation is complete; click **Done**.

3.2.2.2 Configuring and Zoning the LUNs for DDN Devices

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

3.2.3 Configuring Remote Logging of DDN Messages

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

Procedure 13. Configuring the syslog on DDN storage devices

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

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

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

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

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

Note: Use the password that you set in [Configuring the Boot RAID Password When Using DDN Devices](#) on page 50.

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

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

```
CAB01-L1: network
...

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

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

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

Verify the new settings.

```
CAB01-L1: network
...

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

2. Log off the DDN device.

```
CAB01-L1: logout
```

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

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

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

The procedures in this section presuppose that you are familiar with using the SANtricity interface.

3.3.1 Installing SANtricity Storage Manager Software for NetApp, Inc. Engenio Devices

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

Procedure 14. Installing the SANtricity software

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

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

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

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

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

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

3. Set the `DISPLAY` environment variable.

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

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

```
smw:~ # xterm
```

Or

```
smw:~ # xlogo
```

Then, exit the **Xlogo** window or `xterm`.

5. Invoke the executable file.

If you are installing from the CD:

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

Or, if you are installing from a directory:

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

6. Click **Next**. The **License Agreement** window displays.
7. Accept the license agreement and click **Next**. The **Select Installation Type** window displays.
8. Click **Typical (Full Installation)**, then click **Next**.

The **Multi-Pathing Driver Warning** window displays.

9. Click **OK**. The **Pre-Installation Summary** window displays.
10. Click **Install**.

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

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

12. In order to execute `SMclient` from the `crayadm` account, change the ownership and permissions for executable files; otherwise, you must 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
```

3.3.2 Configuring the LUNs for NetApp, Inc. Engenio Devices

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

Procedure 15. Creating the boot RAID volume group for NetApp, Inc. Engenio devices

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

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

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, skip to [step 3](#):
 - **Automatic** — Select this option if you did not assign IP addresses to the storage array controllers using a serial connection. The SANtricity software automatically detects the available controllers, in-band, using the Fibre Channel link.
 - **Manual** — Select this option if you have already assigned IP addresses to the storage array controllers.

Note: The rest of this procedure assumes that you selected the **Manual** option.

3. Double-click the name for the Storage Array that you want to configure. The **Array Management** window displays.
4. Click the **Logical/Physical** tab.

5. Right-click **Unconfigured Capacity** and select **Create Volume**. The **Create Volume** wizard displays.
6. Click **Next** on the **Introduction (Create Volume)** window.
7. Select the **Manual** option on the **Specify Volume Group (Create Volume)** window.
8. Select tray 85, slots 1-4 and click **Add**.
9. Verify that the RAID level is set to 5.
10. Click **Calculate Capacity**.
11. Click **Next** on the **Specify Volume Group (Create Volume)** window.

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

Procedure 16. Creating and configuring volumes for NetApp, Inc. Engenio devices

Configure the boot RAID with at least six LUNs to support the various system management file systems. (The example in this guide configures nine LUNs.)

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

1. Enter a new volume capacity.
2. Specify units as GB or MB.
3. Enter a name.
4. Select the **Customize Settings** option.
5. Click **Next** in the **Specify Capacity/Name (Create Volume)** window.
6. 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.
7. Click **Next** in the **Customize Advanced Volume Parameters (Create Volume)** window.
8. In the **Specify Volume to LUN Mapping** window, select the **Default mapping** option.
9. For **Host type**, select **Linux** from the drop-down menu.

10. Click **Finish** in the **Specify Volume to LUN Mapping** window.
11. When prompted to create more LUNs in the **Creation Successful (Create Volume)** window, select **Yes** unless this is the last volume you are creating. If this is the last volume, select **No** and skip to [step 15](#).
12. In the **Allocate Capacity (Create Volume)** window, verify that **Free Capacity** is selected on **Volume Group 1 (RAID 5)**.
13. Click **Next** in the **Allocate Capacity (Create Volume)** window.
14. Repeat [step 1](#) through [step 13](#) to create all of the volumes described in the boot LUN configuration table in *Installing and Configuring Cray Linux Environment (CLE) Software* (S-2444), which is provided with your CLE release package.
15. Click **OK** in the **Completed (Create Volume)** window.
16. 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. Click **OK**.
 - d. Click **Close**.
17. Exit the tool.

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

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

3.4 Zoning the QLogic FC Switch

If you have a QLogic Fibre Channel Switch, follow [Procedure 17 on page 64](#) to zone the LUNs on your QLogic SANBox switch by using a utility called QuickTools.

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

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

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

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

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

This procedure presupposes that the SANBox is configured and on the HSS network.

Procedure 17. Configuring zoning for a QLogic SANbox switch using QuickTools utility

1. Start a web browser.
2. Enter the IP address of your switch. If your configuration has a single switch, the IP address is 10.1.0.250. The IP address of each RAID controller is preconfigured by Cray and is listed on a sticker on the back of the RAID controller.
3. Enter the login name and password when the **Add a New Fabric** window pops up and prompts for them. The default administrative login name is `admin`, and the default password is `password`.
4. The QuickTools utility displays in your browser. Click **Add Fabric**.

Note: If you receive a dialog box notification that the request failed to connect over a secured connection, click **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, click the **Configured Zonesets** tab.
7. From the toolbar menu, select **Zoning** and then **Edit Zoning**. The **Edit Zoning** window displays.
8. Click 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, click 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.
 - Using the mouse, left-click **Port # 0** and drag it to the BOOT zone.
 - Using the mouse, left-click **Port # 4** and drag it to the BOOT zone. This port is for the boot node.
 - Using the mouse, left-click **Port # 5** and drag it to the BOOT zone. This port is for the SDB node.
 - Using the mouse, left-click **Port # 10** and drag it to the BOOT zone. This port is for the SMW.
13. Click **Apply**. The **error-checking** window displays.
14. When prompted, select **Perform Error Check**.
15. After confirming that no errors were found, click **Save Zoning**.
16. When prompted to activate a Zone Set, click **Yes** and then select the appropriate **XT0** zone set.
17. At this point, Cray recommends that you create a backup of your switch configuration ([Procedure 18 on page 65](#)) before you close and exit the application.

Procedure 18. Creating a backup of your QLogic switch configuration

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

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

1. Start a web browser.
2. Enter the IP address of your switch. The IP address of each RAID controller is preconfigured by Cray and is listed on a sticker on the back of the RAID controller.
3. Enter the login name and password when the **Add a New Fabric** window pops up and prompts for them. The RAID default administrative login name is `admin`, and the default password is `password`.
4. The QuickTools utility appears. Click **Add Fabric**.

Note: If you receive a dialog box that states that the request failed to connect over a secured connection, click **Yes** and continue.

5. From within the QuickTools utility, complete the configuration backup.
 - a. At the top bar, select **Switch** and then **Archive**. A **Save** window pops up with blanks for **Save in:** and **File Name:**.
 - b. Enter the directory (for example, `crayadm`) and a file name (for example, `sanbox_archive`) for saving your QLogic switch configuration.
 - c. Click the **Save** button.
6. Close and exit the application.

3.5 Rediscovering the LUNs

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

Procedure 19. Rebooting the SMW and verifying LUNs are recognized

1. Log on as the `root` user.

```
crayadm@smw:~> su - root
```
2. Enter the following command to ensure that the LUNs are recognized:

```
smw:~ # reboot
```



Warning: Failure to reboot the SMW at this point could produce unexpected results.

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

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

3.6 Partitioning the LUNs

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

Updating or Upgrading Your Cray SMW Software [4]

Cray provides periodic updates to each System Management Workstation (SMW) release. Follow the procedures in this chapter to install an SMW 7.0.UPxx package. The procedures provided in this chapter do **not** change the base operating system version running on your SMW.



Warning: You must be running the SUSE Linux Enterprise Server version 11 Service Pack 2 (SLES 11 SP2) SMW base operating system and a release of SMW 6.0 or SMW 7.0 on your SMW in order to perform the procedures in this chapter.

Note: In the following sections, some examples are left-justified to better fit the page. Left-justification has no special significance.

4.1 Before You Start

- To determine which SLES version is running on your SMW, execute the following command:

```
crayadm@smw:~> cat /etc/SuSE-release
```
- To determine whether there are any changes to the procedures in this chapter, read the *SMW README* and *SMW Errata* that are included in the SMW update directory and read any Field Notices (FNs) related to kernel security fixes.
- For a rack-mount SMW: (Optional) To enable an Integrated Dell Remote Access Controller 6 (iDRAC6), see [Appendix I, Enabling an Integrated Dell Remote Access Controller \(iDRAC6\) on a Rack-mount SMW on page 121](#).

4.2 Backing Up Your Current Software

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

- If you have a rack-mount SMW, see [Appendix E, Rack-mount SMW: Creating a Bootable Backup Drive on page 91](#).
- If you have a desk-side SMW, see [Appendix G, Desk-side SMW: Creating a Bootable Backup Drive on page 107](#).

4.3 Installing an SMW Update Package



Warning: If you are updating a system set that is running, you **must** shut down your Cray system before installing the update package.

Procedure 20. Shutting down the Cray system

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

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

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

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

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

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

Procedure 21. Updating the SMW software and configuration

1. Open a terminal window, and `su` to `root`.

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

2. Mount the release media by using one of the following commands, depending on your media type.

If you are installing the update package from disk, place the Cray SMW 7.0.UPxx Software DVD in the CD/DVD drive and mount it.

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

Note: If you want to install the update package from disk images instead of from the DVD, copy the files to a directory such as `/tmp/SMW_version` on the SMW and then substitute this path for `/media/cdrom` in subsequent instructions.

Or

If you are installing the update package using the ISO image, in the current directory execute the following command with the file name of the `smw-image` ISO image for the update that you are installing. For example:

```
smw:~ # mount -o loop,ro smw-image-7.0.0-1.0700.nnnnn.nnn-n.iso /media/cdrom
```

3. If you have configured postfix on your SMW, complete this step; otherwise, skip to [step 4](#).

To prevent your `master.cf` and `main.cf` postfix configuration files from being recreated during software updates or fixes, ensure the following setting in your `/etc/sysconfig/mail` file on your SMW is set to "no":

```
MAIL_CREATE_CONFIG="no"
```

4. (Optional) If you would like to see what SMW software will be updated in this new release, you may execute these commands prior to your 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
```

Note: Due to the change made that requires the addition of the network variable in the SMWinstall.conf file, these commands will **not** work during an upgrade from SMW 6.0.UPxx to SMW 7.0.UPxx.

5. Create a new copy of the SMWinstall.conf configuration file and modify your new copy of the SMWinstall.conf file for your site-specific requirements. You must be root to modify the SMWinstall.conf configuration file. The SMWinstall.conf configuration file is created during the installation process by copying the SMWinstall.conf template from the distribution media. By default, the SMW configuration file is placed in /home/crayadm/SMWinstall.conf.

Important: If you are upgrading from the SMW 6.0.UPxx release, note the following SMWinstall.conf file changes that are implemented beginning with the SMW 7.0 release:

- In the SMWinstall.conf file, you **must** specify the system interconnection network type for your system; there is no default setting.
- The xttrim utility is used to manage logs on the SMW. The xttrim utility replaced the xtclean_logs utility, so the ENABLE_XTCLEAN_LOGS variable in the SMWinstall.conf file must be set to no. After your SMW software installation is complete, review the xttrim.conf configuration file to ensure the settings are consistent with your site's log maintenance policies.
- The Cray Lightweight Log Manager (LLM) is the log system (LLM=yes in your SMWinstall.conf configuration file). By default, the SMWinstall program enables LLM. The SMWinstall program automatically updates the llm.conf configuration file during the SMW software installation or upgrade. If you make changes to the llm.conf file provided with the release package, you must reapply your changes after an SMW software installation or upgrade. Your llm.conf configuration file will be saved as /etc/opt/cray/llm/llm.conf.rpm save as part of the update process.

- The Linux `sar` service on the SMW is enabled by default.

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 your site has an `/opt/cray/hss/default/etc/xtdiscover.ini` file, the SMW update process does **not** overwrite an existing `xtdiscover.ini` file; the new version is named `xtdiscover.ini.dist`. Compare the content of the new `xtdiscover.ini.dist` with your `xtdiscover.ini` file, and update your `xtdiscover.ini` file appropriately.

Note: If the `xtdiscover.ini` file does **not** exist, then the `xtdiscover.ini.dist` file is copied to the `xtdiscover.ini` file.

8. **If you are upgrading from the from SMW 6.0.UPxx release, complete this step;** otherwise, skip to [step 9](#).

If you customized any file in `/opt/cray/etc`, compare the content of your customized files (for example, `auto.hostname`) in your `/opt/cray/hss/cray_6.0.0/etc` directory to the SMW 7.0.UPxx `/opt/cray/hss/default/etc` files provided with this update (if any). If changes in your customized files must be added to the SMW 7.0.UPxx files provided with this update, make a backup copy of the SMW 7.0.UPxx `/opt/cray/hss/default/etc` files provided with this update before you modify your SMW 7.0.UPxx `/opt/cray/hss/default/etc` files. If necessary, consult your Cray representative before you update your SMW 7.0.UPxx `/opt/cray/hss/default/etc` files.

9. Unmount the SMW 7.0.UPxx Software media.

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

10. If you are using the update disk, eject the SMW 7.0.UPxx Software DVD.

```
smw:~ # eject
```

11. Reboot the SMW.

```
smw:~ # reboot
```

If you have a Cray XE or Cray XK system, continue to [For Cray XE or Cray XK Systems Only: Updating the L0 and L1 Firmware on page 74](#).

4.4 For Cray Cascade Systems Only: Updating the BC and CC Firmware

Cray Cascade 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.

Procedure 22. For Cray Cascade systems only: Updating the BC and CC firmware

1. Update the controller boot image.

```
smw:~ # cat /opt/cray/hss/default/etc/crms-release
HSS-CRMS Tue Oct 23 01:31:09 CDT 2012 on hssbld0 by bwdev
lsb-cray-hss-7.0.0-1.0700.29435.242
smw:~ # hss_make_default_initrd /opt/cray/hss-images/master/7.0.0-1.0700.29435.242
::: Verifying base RPM list to the manifest
::: Installing filesystem hierarchy
::: Installing busybox
::: Installing base RPMs
::: Installing ssh
::: Removing /etc/securetty installed by the pam RPM
::: Installing sshfs
::: Installing rsh
::: Modifying /etc/pam.d/rlogin to remove securetty checking
::: Modifying /etc/pam.d/rsh to remove rhosts and nologin checking (Bug #779466)
::: Installing rsync
::: Installing lsof
::: Installing atftp
::: Installing tcpdump
::: Installing ethtool
::: Installing syslog-ng
::: Installing logrotate
::: Installing ntp/ntpd
::: Installing strace
::: Installing screen
::: Installing minicom
::: Installing ppp
::: Installing mtd-utils
::: Installing /init
::: Installing file.rpm
::: Installing libgmodule
::: Installing Midnight Commander
::: Installing cray-viper
::: Installing spread
::: Installing coreboot-utils
```

```
::: Clearing init.d to be replaced by cray-hss32-filesystem
::: Creating initial etc files needed for root creation
::: Installing Cray kernel
::: Installing latest Cray kernel modules
::: Clearing select /boot items
::: Installing boot-parameters
::: Installing cray-hss32-scripts
::: Installing lsb-cray-hss-controllers
::: Installing cray-libconfig
::: Installing cray-bdm
::: Installing cray-play_xsvf
::: Removing unwanted files from the root
```

=====

The new initrd hierarchy is now in /opt/cray/hss-images/master/7.0.0-1.0700.29435.242.

```
Running hssclone.
Image Clone Complete: /opt/cray/hss-images/default
Running hsspackage.
copying image
copying modules
running depmod
creating load file: /opt/cray/hss-images/default/HSS32.load
compressing initrd.img
Creating pxelinux.0 symlink
Running hssbootlink.
linking /opt/cray/hss-images/default/HSS32/bzImage-3.0.13-0.27.1_1.0700.6694-cray_hss32 \
/opt/tftpboot/bzImage
linking /opt/cray/hss-images/default/HSS32/parameters \
/opt/tftpboot/pxelinux.cfg/default
linking /opt/cray/hss-images/default/HSS32/initrd.img /opt/tftpboot/initrd.img
smw:~ # exit
```

2. Power down the system.

```
crayadm@smw:~> xtcli power down s0
```

3. For each cabinet controller, execute the following command, and then ensure the cabinet controller is up:

```
crayadm@smw:~> xtlogin c0-0
(CC) c0-0 / # /sbin/reboot
crayadm@smw:~> xtalive -l cc c0-0
```

4. Power up the system.

```
crayadm@smw:~> xtcli power up s0
```


5. Update the firmware. Execute the `xtzap` command to update the components.



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

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

Note: If partitions are enabled and `p0` is not, then you must specify each active partition by using `pN`, where *N* designates the partition name:

```
crayadm@smw:~> xtzap -r -v p1
```

6. 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) `chia_fpga` (CHIA FPGA)

This is a list of all blade level components:

`cbb_mc` (CBB BC Microcontroller)

`ibb_mc` (IBB BC Microcontroller)

`anc_mc` (ANC BC Microcontroller)

`bc_bios` (BC Tolapai BIOS)

`lod_fpga` (LOD FPGA)

`node_bios` (Node BIOS)

`loc_fpga` (LOC FPGA)

`qloc_fpga` (QLOC FPGA)

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

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

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

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

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

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

7. Execute the `xtzap -r -v s0` (or `pN`) command again; all firmware revisions should report correctly, except `node_bios`; `node_bios` will display as "NOT_FOUND" until after you complete [step 4.e](#) (the `xtbounce --linktune` command completes).

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

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

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

Continue with [Confirming the SMW is Communicating with the System Hardware on page 75](#).

4.5 For Cray XE or Cray XK Systems Only: Updating the L0 and L1 Firmware



Caution: Your Cray system must be shut down (not powered off) before you begin the next procedure.

Procedure 23. Updating the L0 and L1 firmware

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

Note: The `xtflash` command may take a long time to run, depending on the size of your system.

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

If all L0 daemons do not respond, ensure that your cable connections are secure or contact your 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 your 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.

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

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

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

4.6 Confirming the SMW is Communicating with the System Hardware

Procedure 24. Confirming the SMW is communicating with system hardware

1. Execute the `xtcli status` and 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. For sample output of the `xtcli status` command, see [Confirming the SMW is Communicating with the System Hardware on page 46](#).

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

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

```
crayadm@smw:~> xtmcinfo -t -u s0
```

For sample output of the `xtmcinfo -t -u` command, see [Confirming the SMW is Communicating with the System Hardware on page 46](#).

3. 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
```

If you have a Cray Cascade system, continue with [Updating SMW Software on the Boot Root and Shared Root on page 77](#).

4.7 For Cray XE or Cray XK Systems Only: Verifying Hardware Programmable Interface Controllers (PICs)

Procedure 25. Verifying hardware Programmable Interface Controllers (PICs)

- Execute the `xtcheckpic` command to verify the PICs are current. If any PICs need to be updated, contact your local Cray service personnel. The PIC codes will be updated according to the site staff's priority; the PICs need not be updated while performing this SMW update. However, 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
```

4.8 For Cray Systems with Cray XK Blades Only: Updating the Cray XK Accelerator Firmware

Important: For Cray systems with Cray XK blades: After installing the SMW software update package, you must contact your Cray service personnel to update the Cray XK accelerator firmware so that it is synchronized with the CLE release of the Cray XK driver. These steps **must** be done **before** installing CLE or booting the Cray system and can be run even if your site does not install the latest CLE release package.

4.9 Updating SMW Software on the Boot Root and Shared Root

The `SMWinstallCLE` script updates the SMW software on the boot root and shared root for systems already running the Cray Linux Environment (CLE) software. For more information about the `SMWinstallCLE` script, see the `SMWinstallCLE(8)` man page.

Procedure 26. Updating SMW software on the boot root and shared root

Note: The RPMs that are installed on the boot root and shared root by `SMWinstallCLE` will also be installed the next time that `CLEinstall` is executed if the RPMs were not installed by `SMWinstallCLE`. Use `SMWinstallCLE` if you want to boot CLE after the SMW update. However, if you will be updating CLE software without booting CLE after the SMW update, then this procedure may be skipped.

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

If you are installing the update package from disk, place the `SMW 7.0.UPxx Software DVD` in the CD/DVD drive and mount it.

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

Note: If you want to install the update package from disk images instead of from the DVD, copy the files to a directory such as `/tmp/SMW_version` on the SMW and then substitute this path for `/media/cdrom` in subsequent instructions.

Or

If you are installing the update package using the ISO image, in the current directory execute the following command with the file name of the `smw-image` ISO image for the update that you are installing. For example:

```
smw:~ # mount -o loop,ro smw-image-7.0.0-1.0700.nnnnnn.nnn-n.iso /media/cdrom
```

2. This procedure presupposes that the label BLUE is described in the `/etc/sysset.conf` system set configuration file. The following example updates the BLUE system set from the `/etc/sysset.conf` system set configuration file. See the `sysset.conf(5)` man page for additional information about the `/etc/sysset.conf` file.

```
smw:~ # /media/cdrom/utils/SMWinstallCLE --label=BLUE
```

Note: The `SMWinstallCLE` command checks whether the boot node is booted. If it is booted, `SMWinstallCLE` prompts you to confirm that the system set you are changing is **not** the one booted.

```
HH:MM:SS WARNING: Your bootnode is booted. Please confirm that the system set you
intend to update is not booted.
Do you wish to proceed? [n] y
```

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

3. Unmount the SMW 7.0.UPxx Software media.

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

4. If you are using the update disk, eject the SMW 7.0.UPxx Software DVD.

```
smw:~ # smw:~ # eject
```

Your SMW software is now updated. If you have yet to configure your firewall, see [Appendix D, Setting Up the SUSE Firewall and IP Tables on page 89](#). Then continue to install your CLE software using *Installing and Configuring Cray Linux Environment (CLE) Software (S-2444)*, which is provided with your CLE release package.

Note: To reconfigure a rack-mount SMW's LOGDISK or DBDISK when you need to replace a failed drive with new drive, see [Appendix K, Rack-mount SMW: Replacing a Failed LOGDISK or DBDISK Disk Drive on page 127](#).

Installing RPMs [A]

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.

A.1 Generic RPM Usage

To install RPMs on the SMW, as `root`:

```
smw:~ # rpm -ivh /directorypath/filename.rpm
```


Updating the Time Zone [B]

Procedure 27. Changing the time zone for the SMW and the blade and cabinet controllers



Warning: Perform this procedure while the Cray system is shut down; do not flash blade and cabinet controllers while the Cray system is booted.

You must be logged on as `root`. In this example, the time zone is changed from "America/Chicago" to "America/New_York".

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

```
smw:~ # xtalive -a l0sysd s0
```

2. 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
```

3. Verify that the `zone.tab` file in the `/usr/share/zoneinfo` directory contains the time zone you want to set.

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

4. 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
```

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

```
smw:/etc/sysconfig # 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"
DEFAULT_TIMEZONE="US/Eastern"
```

6. Copy the `/etc/localtime` directory to `/opt/tftpboot` and restart `rsms`.

```
smw:~ # cp /etc/localtime /opt/tftpboot
smw:~ # /etc/init.d/rsms restart
```

For Cray Cascade systems, continue with [step 9](#).

7. **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 8](#).

- 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 10](#).

8. **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 7](#).

- 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 10](#).

9. **For Cray Cascade systems only:** Reboot the cabinet controllers to get the updated time zone.

- a. Power down the system.

```
crayadm@smw:~> xtcli power down s0
```

- b. For each cabinet controller, execute the following command, and then ensure the cabinet controller is up:

```
crayadm@smw:~> xtlogin c0-0
(CC) c0-0 / # /sbin/reboot
crayadm@smw:~> xtalive -l cc c0-0
```

- c. Power up the system.

```
crayadm@smw:~> xtcli power up s0
```

10. Bounce the system.

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


Remote Access to the SMW [C]

Virtual Network Computing (VNC) software enables you to view and interact with the SMW from another computer. The Cray system provides a VNC server, Xvnc; you must download a VNC client to connect to it. See RealVNC (<http://www.realvnc.com/>) or TightVNC (<http://www.tightvnc.com/>) for more information.

Note: The VNC software requires a TCP/IP connection between the server and the viewer. Some firewalls and site security do not allow this connection.

Cray supplies a VNC account `cray-vnc`.

Procedure 28. Starting the VNC server

1. Log on to the SMW as `root` user.
2. Use the `chkconfig` command to check the current status of the server:

```
smw:~ # chkconfig vnc
vnc    off
```

3. Disable `xinetd` startup of Xvnc.

If the `chkconfig` command you executed in [step 2](#) reports that Xvnc was started by INET services (`xinetd`):

```
smw:~ # chkconfig vnc
vnc    xinetd
```

Execute the following commands to disable `xinetd` startup of Xvnc (`xinetd` startup of Xvnc is the SLES 11 default, but it usually is disabled by `chkconfig`):

```
smw:~ # chkconfig vnc off
smw:~ # /etc/init.d/xinetd reload
Reload INET services (xinetd).                                done
```

If no other `xinetd` services have been enabled, the `reload` command will return `failed` instead of `done`. If the `reload` command returns `failed`, this is normal and you can ignore the `failed` notification.

4. Use the `chkconfig` command to start Xvnc at boot time:

```
smw:~ # chkconfig vnc on
```

5. Start the Xvnc server immediately:

```
smw:~ # /etc/init.d/vnc start
```

If the password for `cray-vnc` has not already been established, the system prompts you for one. You must enter a password to access the server.

```
Password: *****
Verify:
Would you like to enter a view-only password (y/n)? n
xauth: creating new authority file /home/cray-vnc/.Xauthority

New 'X' desktop is smw-xt: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-xt:1.log

smw:~ # ps -eda | grep vnc
1839 pts/0    00:00:00 Xvnc
```

Note: The startup script starts the Xvnc server for display :1.

To access the Xvnc server, use a VNC client, such as `vncviewer`, `tight_VNC`, `vnc4`, or a web browser. Direct it to the SMW that is running Xvnc. Many clients allow you to specify whether you want to connect in view-only or in an active mode. If you choose active participation, every mouse movement and keystroke made in your client is sent to the server. If more than one client is active at the same time, your typing and mouse movements are intermixed.

Note: 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.

Procedure 29. For workstation or laptop running Linux: Connecting to the VNC server through an ssh tunnel, using the `vncviewer -via` option

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

Verify that you have the `vncviewer -via` option available. If you do not, use [Procedure 30 on page 87](#).

- If you are connecting from a workstation or laptop running Linux, enter the `vncviewer` command shown below.

The first password you enter is for `crayadm` on the SMW. The second password you enter is for the VNC server on the SMW, which was set when the VNC server was started for the first time using `/etc/init.d/vnc start` on the SMW.

```
/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
```

Procedure 30. For workstation or laptop running Linux: Connecting to the VNC server through an ssh tunnel

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

1. This `ssh` command starts an `ssh` session between the local Linux computer and the SMW, and it also creates an SSH tunnel so that port 5902 on the localhost is forwarded through the encrypted SSH tunnel to port 5901 on the SMW. You will be prompted for the `crayadm` password on the SMW.

```
local_linux_prompt> ssh -L 5902:localhost:5901 smw -l crayadm
Password:
crayadm@smw>
```

2. Now `vncviewer` can be started using the local side of the SSH tunnel, which is port 5902. You will be prompted for the password of the VNC server on the SMW. This password was set when the VNC server was started for the first time using `/etc/init.d/vnc start` on the SMW.

```
local_linux_prompt> vncviewer localhost:2
Connected to RFB server, using protocol version 3.7
Performing standard VNC authentication
Password:
```

The VNC window from the SMW appears. All traffic between the `vncviewer` on the local Linux computer and the VNC server on the SMW is now encrypted through the SSH tunnel.

Procedure 31. For workstation or laptop running Mac OS X: Connecting to the VNC server through an ssh tunnel

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

1. This `ssh` command starts an `ssh` session between the local Mac OS X computer and the SMW, and it also creates an SSH tunnel so that port 5902 on the localhost is forwarded through the encrypted SSH tunnel to port 5901 on the SMW. You will be prompted for the `crayadm` password on the SMW.

```
local_mac_prompt> ssh -L 5902:localhost:5901 smw -l crayadm
Password:
crayadm@smw>
```

2. Now `vncviewer` can be started using the local side of the SSH tunnel, which is port 5902. You will be prompted for the password of the VNC server on the SMW. This password was set when the VNC server was started for the first time using `/etc/init.d/vnc start` on the SMW.

If you type this on the Mac OS X command line after having prepared the SSH tunnel, the `vncviewer` will pop up:

```
local_mac_prompt% open vnc://localhost:5902
```

The VNC window from the SMW appears. All traffic between the `vncviewer` on the local Mac OS X computer and the VNC server on the SMW is now encrypted through the SSH tunnel.

Procedure 32. For workstation or laptop running Windows: Connecting to the VNC server through an `ssh` tunnel

Note: If you are connecting from a computer running Windows, then both a VNC client program, such as TightVNC and an SSH program, such as PuTTY, SecureCRT, or OpenSSH are recommended.

1. The same method described in [Procedure 30](#) 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: Details about how to create the SSH tunnel vary amongst the different SSH programs for Windows computers.

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 click on the **Connect** button.

Setting Up the SUSE Firewall and IP Tables [D]

Procedure 33. Setting up the SUSE firewall and IP tables

Note: You can implement this procedure on a running system.

1. By using the `SuSEfirewall2` program and the following procedure, change the IP tables rules to close off all unnecessary ports on the SMW.

```
smw:~ # iptables -L
smw:~ # vi /etc/sysconfig/SuSEfirewall2
```

Change the settings of the following variables to be as shown below:

```
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/SuSEfirewall2.orig /etc/sysconfig/SuSEfirewall2
99c99
< FW_DEV_EXT="eth-id-00:30:48:5c:b0:ee eth0"
---
> FW_DEV_EXT="any eth0"
114c114
< FW_DEV_INT="eth-id-00:0e:0c:b4:df:64 eth-id-00:0e:0c:b4:df:65
    eth-id-00:0e:0c:b4:df:66 eth-id-00:0e:0c:b4:df:67 eth1 eth2 eth3 eth4"
---
> FW_DEV_INT="eth1 eth2 eth3 eth4 lo"
263c263
< FW_SERVICES_EXT_UDP=" "
---
> FW_SERVICES_EXT_UDP="161"
394c394
< FW_TRUSTED_NETS=" "
---
> FW_TRUSTED_NETS="10.3.1.254,tcp,7004 10.5.1.2,udp,514 10.5.1.2,tcp,6811:6815"
```

2. Invoke the modified configuration.

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

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

```
smw:~ # chkconfig SuSEfirewall12_init on  
smw:~ # chkconfig SuSEfirewall12_setup on
```

4. To verify your changes to the `iptables`, issue the following command:

```
smw:~ # iptables -L
```

Rack-mount SMW: Creating a Bootable Backup Drive [E]

This appendix describes the procedure to create a System Management Workstation (SMW) bootable backup drive for a rack-mount SMW. For information about how to modify a bootable backup drive in order to boot from and run the SMW from the backup root partition, see [Appendix F, Rack-mount SMW: Setting up the Bootable Backup Drive as an Alternate Boot Device](#) on page 103.

E.1 Rack-mount SMW: Creating a Cray System Management Workstation (SMW) Bootable Backup Drive



Warning: You **must** be running the SUSE Linux Enterprise Server version 11 Service Pack 2 (SLES 11 SP2) SMW base operating system and a release of SMW 6.0 or SMW 7.0 on your SMW in order to perform the procedures in this chapter.

The following procedure creates a bootable backup drive for a rack-mount SMW in order 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.

Procedure 34. Rack-mount SMW: Creating an SMW bootable backup drive

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



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

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

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

```
crayadm@smw:~> su -
Password:
smw:~ # ls -al /dev/disk/by-path
total 0
drwxr-xr-x 2 root root 380 Mar 15 13:21 .
drwxr-xr-x 6 root root 120 Mar 11 18:42 ..
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:00:11.0-scsi-0:0:0:0 -> ../../sr0
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:00:12.2-usb-0:3:1.0-scsi-0:0:0:0 -> ../../sdf
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:00:12.2-usb-0:3:1.1-scsi-0:0:0:0 -> ../../srl
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:00:12.2-usb-0:3:1.1-scsi-0:0:0:1 -> ../../sdg
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:00:13.2-usb-0:2.1:1.0-scsi-0:0:0:0 -> ../../sde
lrwxrwxrwx 1 root root 10 Mar 11 18:42 pci-0000:00:13.2-usb-0:2.1:1.0-scsi-0:0:0:0-part1 -> ../../sdel
lrwxrwxrwx 1 root root 10 Mar 11 18:42 pci-0000:00:13.2-usb-0:2.1:1.0-scsi-0:0:0:0-part2 -> ../../sde2
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:05:00.0-sas-phy4-...
lrwxrwxrwx 1 root root 10 Mar 14 15:57 pci-0000:05:00.0-sas-phy4-...
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:05:00.0-sas-phy5-...
lrwxrwxrwx 1 root root 10 Mar 14 16:00 pci-0000:05:00.0-sas-phy5-...
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:05:00.0-sas-phy6-...
lrwxrwxrwx 1 root root 10 Mar 15 13:21 pci-0000:05:00.0-sas-phy6-...
lrwxrwxrwx 1 root root 10 Mar 15 13:21 pci-0000:05:00.0-sas-phy6-...
lrwxrwxrwx 1 root root 9 Mar 11 18:42 pci-0000:05:00.0-sas-phy7-...
lrwxrwxrwx 1 root root 10 Mar 11 18:42 pci-0000:05:00.0-sas-phy7-...
lrwxrwxrwx 1 root root 10 Mar 11 18:42 pci-0000:05:00.0-sas-phy7-...
```

2. **After you have installed SLES 11 SP2 as your SMW base operating system, complete this step to modify the configuration files on the SMW root drive, which will standardize the SMW's boot-time drive names with the Linux run-time drive names. If your SMW configuration files on the SMW root drive have been modified already (that is, your site has completed this step at least once after installing SLES 11 SP2 as your SMW base operating system, skip to [step 3](#).**

Set up ordered drives on your rack-mount SMW.

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

Provided with the SMW 7.0.UP01 release is a script named `smwmapdrives.sh`. This script will identify local (internal) drives mounted in the SMW and provide their Linux device (`/dev`) names. Execute the `smwmapdrives.sh` script on the SMW.

```

smw:~ # ./smwmapdrives.sh
List of SMW-installed disk drives
-----
Physical slot 0:
    /dev/sda
    /dev/disk/by-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RDS
    /dev/disk/by-id/scsi-SATA_FUJITSU_MHZ2160_K85DTB227RDS
    /dev/disk/by-path/pci-0000:05:00.0-sas-phy7-0x4433221107000000-lun-0
Physical slot 1:
    /dev/sdc
    /dev/disk/by-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RD7
    /dev/disk/by-id/scsi-SATA_FUJITSU_MHZ2160_K85DTB227RD7
    /dev/disk/by-path/pci-0000:05:00.0-sas-phy6-0x4433221106000000-lun-0
Physical slot 2:
    /dev/sdd
    /dev/disk/by-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RF3
    /dev/disk/by-id/scsi-SATA_FUJITSU_MHZ2160_K85DTB227RF3
    /dev/disk/by-path/pci-0000:05:00.0-sas-phy5-0x4433221105000000-lun-0
Physical slot 3:
    /dev/sdb
    /dev/disk/by-id/ata-ST9500620NS_9XF0665V
    /dev/disk/by-id/scsi-SATA_ST9500620NS_9XF0665V
    /dev/disk/by-path/pci-0000:05:00.0-sas-phy4-0x4433221104000000-lun-0
Physical slot 4:
    NOT INSTALLED
Physical slot 5:
    NOT INSTALLED

```

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

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

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

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

- b. Back up the following files before proceeding:

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

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

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

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

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

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

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

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

Note: The output of the `smwmapdrives.sh` script in [step 2.a](#) should be referenced for drive namings.

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

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

For example, if the root and swap file systems are currently configured to mount `/dev/sda2`, they should be changed. Using the `by-id` device name from the example in [step 2.a](#), 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-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RDS-part1 swap swap defaults 0 0
/dev/disk/by-id/ata-FUJITSU_MHZ2160BK_G2_K85DTB227RDS-part2 / ext3 acl,user_xattr 1 1
```

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

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

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

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

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

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

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

In this example, the partition table consists of 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. (Example output spacing was modified to fit on the printed page.)

```
mw:~ # 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: 160.0 GB, \
160041885696 bytes
255 heads, 63 sectors/track, 19457 cylinders, total 312581808 sectors
Units = sectors of 1 * 512 = 512 bytes
Disk identifier: 0x00000082

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

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

```
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: 160.0 GB, \
160041885696 bytes
255 heads, 63 sectors/track, 19457 cylinders, total 312581808 sectors
Units = sectors of 1 * 512 = 512 bytes
Disk identifier: 0x00000080

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

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

Command (m for help): d
Selected partition 1

Command (m for help): n
Command action
```



```

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

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

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

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

```

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

```

smw:~ # fdisk -lu /dev/disk/by-path/pci-0000:05:00.0-sas-phy6-0x4433221106000000-lun-0
Disk /dev/disk/by-path/pci-0000:05:00.0-sas-phy6:1-0x4433221106000000:0-lun0: 160.0 GB, \
160041885696 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 only necessary 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.

Note: 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 only detects the first 16 devices. Therefore, if you add a disk to the SMW **after** the SMW is connected to the boot RAID, it is advisable to reboot the SMW **before** continuing this procedure.

- 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.

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

For example:

```
smw:~ # grub
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)l+17 p (hd2,1)/boot/grub/stage2 \
/boot/grub/menu.lst"... succeeded
Done.
grub> quit
```

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

11. Unmount the backup root partition.

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

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

Rack-mount SMW: Setting up the Bootable Backup Drive as an Alternate Boot Device [F]

The following procedure modifies a bootable backup drive, generated in [Appendix E, Rack-mount SMW: Creating a Bootable Backup Drive on page 91](#), in order to boot from and run the SMW from the backup root partition. For information about how to recover the SMW, see the Cray publication *Managing System Software* (S–2393).

F.1 Rack-mount SMW: Setting Up the Bootable Backup Drive as an Alternate Boot Device



Warning: You **must** be running the SUSE Linux Enterprise Server version 11 Service Pack 2 (SLES 11 SP2) SMW base operating system and a release of SMW 6.0 or SMW 7.0 on your SMW in order to perform the procedures in this chapter.

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

Important: 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.

Procedure 35. Rack-mount SMW: Setting up the bootable backup drive as an alternate boot device

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



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

1. Mount the backup drive's root partition.

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

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

An example `/boot/grub/menu.lst` file follows. 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 SP2 - 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 SP2 - 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 SP2 - 3.0.26-0.7
    root (hd0,1)
    kernel /boot/vmlinuz-3.0.26-0.7-default \
    root=/dev/disk/by-path/pci-0000:05:00.0-sas-phy7-0x4433221107000000-lun-0-part2 \
    showopts ide=nodma apm=off noresume edd=off powersaved=off \
    nohz=off highres=off processor.max_cstate=1 nomodeset x11failsafe vga=0x31a
    initrd /boot/initrd-3.0.26-0.7-default
```

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

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

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

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

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

4. Unmount the backup drive.

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

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

Desk-side SMW: Creating a Bootable Backup Drive [G]

This appendix describes the procedure to create a System Management Workstation (SMW) bootable backup drive for a desk-side SMW. For information about how to modify a bootable backup drive for a desk-side SMW in order to boot from and run the SMW from the backup root partition, see [Appendix H, Desk-side SMW: Setting Up the Bootable Backup Drive as an Alternate Boot Device](#) on page 117.

G.1 Desk-side SMW: Creating an System Management Workstation (SMW) Bootable Backup Drive

The following procedure creates a System Management Workstation (SMW) bootable backup drive for a desk-side SMW in order 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.

Note: In the following procedure, `/dev/sdX2` is the SMW disk (either `/dev/sdb2` or `/dev/sdc2`).

Procedure 36. Desk-side SMW: Creating an SMW bootable backup drive

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



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

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

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

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

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

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

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

- Slice 1: 4 GB Linux swap partition
 - Slice 2: Balance of disk space used for the root file system
- a. Use the `fdisk` command to display the boot disk partition layout.

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

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

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

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

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

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

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

```
Command (m for help): d
Partition number (1-5): 2
Command (m for help): d
Selected partition 1
Command (m for help): n
Command action
```

```

e extended
p primary partition (1-4)
p
Partition number (1-4): 1
First sector (63-625105215, default 63): (Press the Enter key)
Using default value 63
Last sector or +size or +sizeM or +sizeK (63-625105215, default 625105215): 8401994

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

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

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

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

```

c. Display the boot backup disk partition layout.

```

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

```

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

3. Initialize the swap device.

```

smw:~ # mkswap /dev/sdb1

```

4. Standardize the /etc/fstab and grub disk device names.

The device names that the installation process writes into the /boot/grub/menu.lst file are UDEV-based names (for example, /dev/disk/by-id/scsi-SATA_ST3320620AS_922J3-part2 or /dev/disk/by-id/ata-ST3320620A_9QFA85PV-part2) instead of the more commonly used device names (for example, /dev/sda2 or

/dev/hda2). In the following procedures, edit the `/boot/grub/menu.lst` file to change ONLY the long UDEV-based name to the shorter, commonly used device name reflected in the output of the `df` command on your system.

If the device names have already been standardized, skip to [step 5](#).



Caution: Mistakes in the `/boot/grub/menu.lst` file will affect your ability to boot the SMW.

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

```
smw:~ # head -2 /etc/fstab
/dev/disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part2 /      ext3    acl,user_xattr  1 1
/dev/disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part1 swap  swap  defaults        0 0

smw:~ # more /boot/grub/menu.lst
###Don't change this comment - YaST2 identifier: Original name: linux###
title SUSE Linux Enterprise Server 11 - 2.6.27.19-5
    root (hd0,1)
    kernel /boot/vmlinuz-2.6.27.19-5-default \
    root=/dev/disk/by-id/ata-ST3320620AS_5QF00F84-part2 \
    resume=/dev/sda1 splash=silent crashkernel=256M-:128M@16M
    showopts vga=0x31a initrd /boot/initrd-2.6.27.19-5-default

###Don't change this comment - YaST2 identifier: Original name: failsafe###
title Failsafe -- SUSE Linux Enterprise Server 11 - 2.6.27.19-5
    root (hd0,1)
    kernel /boot/vmlinuz-2.6.27.19-5-default \
    root=/dev/disk/by-id/ata-ST3320620AS_5QF00F84-part2 showopts \
    ide=nodma apm=off noresume edd=off powersaved=off nohz=off highres=off
    processor.max_cstate=1 xllfailsafe vga=0x31a
    initrd /boot/initrd-2.6.27.19-5-default
```

- b. Execute the `df` command to get the name of the device to use in the `/etc/fstab` and `/boot/grub/menu.lst` files to replace the long UDEV-based device name. Then edit your `/etc/fstab` and `/boot/grub/menu.lst` files appropriately.
 - 1) 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
```

- 2) Save a copy of your `/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
```

- 3) Edit your `/etc/fstab` file appropriately; you will use the device name (`dev`) you got from the `df` command output. In

this example, the "1" and "2" refer to the partition names on the device. Change the following lines, which changes the long name `disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part2` to `sda2` and changes `disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part1` to `sda1`. Ensure that your swap is on `sda1`:

```
/dev/disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part2 /      ext3  acl,user_xattr  1 1
/dev/disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part1 swap swap  defaults          0 0
```

to:

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

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

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

to:

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

and change the following lines:

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

to:

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

- 5) Verify that the edited files are correct and matches the output of the `df` command.

```
smw:~ # head -2 /etc/fstab
/dev/sda2  /      ext3    acl,user_xattr  1 1
/dev/sda1  swap   swap    defaults        0 0

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

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

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



Caution: Although all of the disks connected to the SMW are available to the system, grub only detects the first 16 devices. Therefore, if you add a disk to the SMW **after** the SMW is connected to the boot RAID, it is advisable to reboot the SMW **before** continuing 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, utilizing the standardized drive naming. This file can be viewed for verification; for example:

```
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. Install the grub boot loader.

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



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

- 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_SDX
```

- 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.

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

For example:

```
smw:~ # grub
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
```

11. Unmount the backup root partition.

```
smw:~ # umount /dev/sdb2
```

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

Desk-side SMW: Setting Up the Bootable Backup Drive as an Alternate Boot Device [H]

The following procedure modifies a bootable backup drive, generated in [Appendix G, Desk-side SMW: Creating a Bootable Backup Drive on page 107](#), in order to boot from and run the SMW from the backup root partition. For information about how to recover the SMW, see the Cray publication *Managing System Software* (S-2393).

H.1 Desk-side SMW: Setting Up the Bootable Backup Drive as an Alternate Boot Device

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

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

Procedure 37. Desk-side SMW: Setting up the bootable backup drive as an alternate boot device

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



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

1. Mount the backup drive's root partition.

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

2. Create a new boot entry in the `/boot/grub/menu.lst` file. This entry should be a duplicate of the primary boot entry with the following changes:
 - Modify the title to uniquely identify the backup boot entry.

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

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

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

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

###Don't change this comment - YaST2 identifier: Original name: failsafe###
title Failsafe -- SUSE Linux Enterprise Server 11 - 2.6.27.19-5
  root (hd0,1)
  kernel /boot/vmlinuz-2.6.27.19-5-default root=/dev/sda2 showopts \
  ide=nodma apm=off noresume edd=off powersaved=off nohz=off highres=off \
  processor.max_cstate=1 xllfailsafe vga=0x31a \
  initrd /boot/initrd-2.6.27.19-5-default

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

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

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

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

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

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

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

4. Unmount the backup drive.

```
smw:~ # umount /dev/sdx2
```

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

Enabling an Integrated Dell Remote Access Controller (iDRAC6) on a Rack-mount SMW [I]

Enabling an Integrated Dell Remote Access Controller (iDRAC6) allows you to manage your rack-mount SMW remotely.

I.1 Before You Start

Before you enable an iDRAC6 on a rack-mount SMW, you must:

- Have physical access to the SMW console
- Know your iDRAC6 IP address.
- Know your SMW `root` account password
- Have Java installed on your SMW console

I.2 Enabling an Integrated Dell Remote Access Controller (iDRAC6) on a Rack-mount SMW

Procedure 38. Enabling an Integrated Dell Remote Access Controller (iDRAC6) on a rack-mount SMW

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

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

2. Connect Ethernet cable to the iDRAC6 port. The cable is located on back of a rack-mount SMW in the lower left corner.
3. Power up the SMW.
4. After the BIOS, Dell PowerEdge Expandable RAID Controller (PERC) card, and disk map have displayed, the IPv4/IPv6 information displays. When the IPv4/IPv6 information displays, press `Ctrl-E`.
5. Using the arrow keys, select **Lan Parameters**, then press `Enter`.
6. Select **NIC Selection** and set it to **Dedicated**. Then press `Esc`.

7. Using the arrow keys, scroll down and select the **IPv4 settings** section.
 - a. Ensure that IPv4 is enabled.
 - b. Confirm that the IPv4 address source is set to static:
`IPv4 Address Source: Static`
 - c. Enter your iDRAC6 IP addresses for the following:
 - Address:
 - Subnet Mask:
 - Default Gateway:
 - d. Ensure that IPv6 is disabled.
 - e. Press **Esc** and return to the **Lan Parameters** window.
8. Using the arrow keys, select **Lan User Configuration**, then press **Enter**.

Note: This configuration is for both `ssh` and web browser access to the iDRAC.

 - a. Enter the `root` account name and SMW password:
`Account User name: root`
`Enter Password: *****`
`Confirm Password: *****`
 - b. Press **Esc**.
9. Press **Esc** again.
10. Select **Save Changes and Exit**, then press **Enter**. The SMW will complete booting up; no user intervention is required.

I.3 Using the iDRAC6

Procedure 39. Using the iDRAC6

1. Bring up a web browser.
2. Go to: `https://cray-drac`. A login screen appears.
3. Enter the account user name and password that you set up in [Procedure 38 on page 121, step 8.a](#). Then click on **Submit**.

The **System Summary** window appears.
4. To access the SMW console, click on the **Console Media** tab.

The **Virtual Console and Virtual Media** window appears.
5. Click on **Launch Virtual Console**.

Tip: By default, your console window has two cursors: one for the console and one for your own windowing environment. To switch to single-cursor mode, click on **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 click on **File**, then **Exit**. You will still be logged into the iDRAC6 in your web browser.

For detailed information, see the iDRAC6 documentation at:

<http://support.dell.com/support/edocs/software/smdrac3/idrac>

Rack-mount SMW: Installing Server Administrator Software [J]

For advanced control over the Integrated Dell Remote Access Controller (iDRAC) to provide features such as Automatic Recovery (automatic system boot after a power event), the Server Administrator software must be installed from the Dell Systems Management Tools and Documentation DVD. This DVD was provided with your rack-mount SMW.

Procedure 40. Rack-mount SMW: Installing the Dell Systems Management Tools and Documentation DVD

1. Obtain the Dell System Management Tools and Documentation DVD.
2. Login as `root` on the SMW.
3. Mount the DVD.

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

4. Go to the location of the installation scripts.

```
smw:~ # cd /media/cdrom/SYSMGMT/srvadmin/linux/supportscripts
```

5. Execute the script to install the software.

```
smw:~ # sh srvadmin-install.sh --express
```

6. Start the Server Administrator services.

```
smw:~ # sh srvadmin-services.sh start
```

7. Double click the icon named **Launch Server Administrator** on the SMW screen.
8. Enter the SMW user name `root`.
9. Enter the SMW `root` account password.

Now you are ready to manage the system for Properties, Shutdown, Logs, AlertManagement, and Session Management.

The following websites may be useful:

- Dell OpenManage Linux Repository — See <http://linux.dell.com/wiki/index.php/Repository/OMSA>

- Dell OpenManage Server Administrator Version 6.5 Manuals (including complete installation guide and user guide — See <http://support.dell.com/support/edocs/software/svradmin/7.0/en/index.htm>)

Procedure 41. Rack-mount SMW: Configuring the SMW to restore AC power after a power failure

Note: To display the iDRAC IP information, execute the command:

```
/usr/bin/ipmitool -I open lan print 1
```

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 single click on 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**, single click on the value field, check **Disable**, and then select **Apply Change**.

12. Exit the GUI.

Rack-mount SMW: Replacing a Failed LOGDISK or DBDISK Disk Drive [K]

K.1 Rack-mount SMW: Replacing a Failed LOGDISK or DBDISK Disk Drive



Warning: You **must** be running the SUSE Linux Enterprise Server version 11 Service Pack 2 (SLES 11 SP2) SMW base operating system and a release of SMW 6.0 or SMW 7.0 on your SMW in order to perform the procedures in this chapter.

Procedure 42. Rack-mount SMW: Replacing a failed LOGDISK or DBDISK disk drive

Note: This procedure specifies replacing the LOGDISK disk. If you are replacing the DBDISK, use the appropriate `/dev/disk/by-path/pci*` device name.

1. Replace the failed drive with the new drive.
2. Reboot the SMW.

```
smw:~ # reboot
```

3. Reconfigure LOGDISK.

```
smw:~ # /sbin/fdisk  
/dev/disk/by-path/pci-0000:05:00.0-sas-phy4-0x4433221104000000-lun-0  
Command (m for help): p
```

- a. Delete all the current partitions, if there are any.

```
Command (m for help): d  
Partition number 4  
Command (m for help): d  
Partition number 3  
Command (m for help): d  
Partition number 2  
Command (m for help): d  
Partition number 1  
Command (m for help): p
```

b. Create the new, single partition.

```
Command (m for help): n
Command action
  e   extended
  p   primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-121601, default 1): # Hit return, take the default
Using default value 1
Last cylinder, +cylinders or +size{K,M,G} (1-121601, default 121601): # Hit
return, 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     97676001     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. Restart rsms and mazama daemons.

```
smw:~ # /etc/init.d/rsms restart
smw:~ # /etc/init.d/cray-mzsd restart
smw:~ # /etc/init.d/cray-mzwatcher restart
smw:~ # /etc/init.d/dbMonitor restart
```