



XC™ Series SMW HA Installation Guide (SLEHA12.SP0.UP03) S-0044

Contents

1 About the XC™ Series SMW HA Installation Guide.....	4
1.1 Related Publications.....	7
1.2 Distribution Media.....	8
2 Install and Configure an SMW HA System.....	9
2.1 Prepare for an SMW HA Fresh Install.....	9
2.2 Install and Configure the XC System on the First SMW.....	10
2.3 Prepare to Install SMW HA Software.....	10
2.3.1 Record All Site Customization and Local Changes to the First SMW.....	10
2.3.2 Power Off the Second SMW.....	11
2.3.3 Verify that eth0 and eth3 are Unmanaged Interfaces.....	11
2.4 Install the SMWHA Software on the First SMW.....	12
2.5 Install and Configure the Second SMW	14
2.5.1 Install the Base Operating System on the Second SMW.....	15
2.5.2 Install SMW and CLE Software on the Second SMW.....	41
2.5.3 Configure the Second SMW for CLE System Hardware.....	49
2.5.4 Install the SMWHA Software on the Second SMW.....	51
2.6 Reboot the Second SMW and Power On the First SMW.....	52
2.7 Configure the SMW HA Cluster.....	53
2.7.1 Gather SMW HA Cluster Information.....	53
2.7.2 Configure Required Cluster Settings.....	55
2.7.3 Add the Second SMW to smw_nodes Node Group.....	64
2.8 Change Default HA Passwords After Installation	65
2.8.1 Change the Default iDRAC Password.....	66
2.9 Configure Failover Notification.....	66
2.10 Configure the Power Management Database with DRBD for SMW HA.....	68
2.11 Finish Configuring the SMW HA System.....	72
3 Update an SMW HA System.....	74
3.1 Prepare for an SMW HA Update.....	74
3.1.1 Start a Typescript File on the First/Active SMW.....	75
3.1.2 Record the Cluster Configuration.....	76
3.1.3 Verify that eth0 and eth3 are Unmanaged Interfaces.....	76
3.1.4 Verify Cluster in a Sane State.....	77
3.1.5 Exit Typescript and Shut Down Second SMW.....	78
3.2 Update and Configure the XC System on the First SMW.....	78
3.3 Update the SMWHA Software on the First SMW.....	78

3.4 Update and Configure the XC System on the Second SMW.....	81
3.4.1 Start a Typescript File on the Second/Passive SMW.....	81
3.4.2 Collect Software Media.....	82
3.4.3 Mount Software Media and Prepare install.cle.conf.....	83
3.4.4 Run the Installer for a Software Update.....	84
3.4.5 Set Default Snapshot and Boot the SMW during a Software Update.....	86
3.5 Make a Post-update Snapshot using snaputil.....	87
3.6 Update the SMWHA Software on the Second SMW.....	88
3.7 Reboot the Second SMW after an Update.....	89
3.8 Configure the SMW HA Cluster after an Update.....	89
3.9 Revert to a Previous SMW HA Snapshot.....	93
4 Customize a Preinstalled SMW HA System.....	95
4.1 Change IP Addresses, Host Names, Name Server, and Route Settings on the First SMW.....	96
4.2 Change the Cluster Configuration on the First SMW.....	98
4.3 Change IP Addresses, Host Names, Name Server, and Route Settings on the Second SMW.....	100
4.4 Finish Customizing a Preinstalled SMW HA System.....	102
4.5 Verify Cluster Status After Customization.....	103
4.6 Change Default SMW, iDRAC, and STONITH Passwords After Customization.....	104
4.6.1 Change the Default iDRAC Password.....	105
5 Optional Cluster Configuration Changes.....	106
5.1 Rules for Changing the SMW HA Cluster Configuration.....	106
5.2 Change Failover Notification.....	106
5.3 Add Site-specific Files to the Synchronization List.....	107
5.3.1 Synchronized Files.....	108
5.4 Set the Migration Threshold for a Resource.....	108
6 Verify the SMW HA Cluster Configuration.....	110
7 Additional Procedures for an Installed SMW HA System.....	113
7.1 Migrate PostgreSQL Data to DRBD for an SMW HA System.....	113
7.2 Enable Multipath on an Installed SMW HA System.....	119
7.3 Re-create Host Certificates to Remedy SSL Certificate Verification Failure.....	123

1 About the XC™ Series SMW HA Installation Guide

Scope and Audience

The *XC™ Series SMW HA Installation Guide* (S-0044) includes procedures for installing a Cray XC™ Series system that includes two System Management Workstations (SMW) configured for High Availability (HA), also called SMW failover or an SMW HA cluster. An SMW HA system is a Cray XC system with two second-generation rack-mount SMWs, either Dell R815 or Dell R630 models. The SMWs run the SUSE Linux Enterprise High Availability (SLEHA) Extension and the Cray SMW High Availability Extension release package, also called the *SMW HA package*. The two SMWs must have the same hardware, software, and configuration settings.

This publication provides HA-specific installation and configuration procedures for a full initial installation, a software update, and the customization of a pre-installed SMW HA system when it arrives on the customer site. Only the HA-specific procedures are found in this guide. See *XC™ Series Software Installation and Configuration Guide* for the procedures to install, configure, and customize the first SMW.

This publication does not include administration procedures. For information on managing a running SMW HA system, see the *SMW HA Administration Guide for XC Series Systems* (S-2551).

This publication is intended for system installers, administrators, and anyone who installs and configures SMW HA software on a Cray XC Series system. It assumes some familiarity with standard Linux and open source tools (e.g., zypper/yum for RPMs, Ansible, YAML/JSON configuration data).

SMW HA SLEHA12.SP0.UP03 Release

XC™ Series SMW HA Installation Guide (SLEHA12.SP0.UP03) S-0044 supports Cray software release SMW HA SLEHA12.SP0.UP03 for Cray XC™ Series systems, released on 16 February 2017.

In previous releases, this publication was titled *SMW HA Initial Installation Guide for XC Series Systems* (SLEHA12.SP0.UP01), *SMW HA Installation Guide* (S-0044-F), and *Installing, Configuring, and Managing SMW Failover on the Cray XC System* (S-0044-D).

New in this release

- Workaround to address incorrect zypper repo setup in these two procedures:
 - [Reboot the Second SMW and Power On the First SMW](#) on page 52
 - [Reboot the Second SMW after an Update](#) on page 89
- New procedure to use on an SMW HA system if DataWarp DataWarp `dwcli` and `dwstat` commands result in CERTIFICATE_VERIFY_FAILED errors: [Re-create Host Certificates to Remedy SSL Certificate Verification Failure](#) on page 123
- New procedure: [Revert to a Previous SMW HA Snapshot](#) on page 93
- New step to customize the DRBD configuration file (`/etc/drbd.d/r0.res`) to change the host name and email address for split-brain notification in these two procedures:
 - [Change the Cluster Configuration on the First SMW](#) on page 98

- [Change IP Addresses, Host Names, Name Server, and Route Settings on the Second SMW](#) on page 100
- New SMWHAconfig log file path in these two procedures:
 - [Configure Required Cluster Settings](#) on page 55
 - [Configure the SMW HA Cluster after an Update](#) on page 89
- The `crm_resource --list` command now shows the status of the resources in addition to the resource names.

Host Name Conventions for SMW HA Systems

These host name conventions are used to refer to the SMWs in an HA cluster:

<code>smw1#</code>	Specifies the currently active SMW.
<code>smw2#</code>	Specifies the currently passive SMW.
<code>virtual-smw#</code>	Specifies the virtual (active) SMW, which could be either <code>smw1</code> or <code>smw2</code> .

Command Prompt Conventions

Host name and account in command prompts The host name in a command prompt indicates where the command must be run. The account that must run the command is also indicated in the prompt.

- The `root` or super-user account always has the `#` character at the end of the prompt.
- Any non-`root` account is indicated with `account@hostname>`. A user account that is neither `root` nor `crayadm` is referred to as `user`.

<code>smw#</code>	Run the command on the SMW as <code>root</code> .
<code>cmc#</code>	Run the command on the CMC as <code>root</code> .
<code>sdb#</code>	Run the command on the SDB node as <code>root</code> .
<code>crayadm@boot></code>	Run the command on the boot node as the <code>crayadm</code> user.
<code>user@login></code>	Run the command on any login node as any non- <code>root</code> user.
<code>hostname#</code>	Run the command on the specified system as <code>root</code> .
<code>user@hostname></code>	Run the command on the specified system as any non- <code>root</code> user.

smw1# smw2#	For a system configured with the SMW failover feature there are two SMWs—one in an active role and the other in a passive role. The SMW that is active at the start of a procedure is <i>smw1</i> . The SMW that is passive is <i>smw2</i> .
smwactive# smwpassive#	In some scenarios, the active SMW is <i>smw1</i> at the start of a procedure—then the procedure requires a failover to the other SMW. In this case, the documentation will continue to refer to the formerly active SMW as <i>smw1</i> , even though <i>smw2</i> is now the active SMW. If further clarification is needed in a procedure, the active SMW will be called <i>smwactive</i> and the passive SMW will be called <i>smwpassive</i> .

Command prompt inside chroot If the `chroot` command is used, the prompt changes to indicate that it is inside a chroot environment on the system.

```
smw# chroot /path/to/chroot
chroot-smw#
```

Directory path in command prompt Example prompts do not include the directory path, because long paths can reduce the clarity of examples. Most of the time, the command can be executed from any directory. When it matters which directory the command is invoked within, the `cd` command is used to change into the directory, and the directory is referenced with a period (.) to indicate the current directory.

For example, here are actual prompts as they appear on the system:

```
smw:~ # cd /etc
smw:/etc# cd /var/tmp
smw:/var/tmp# ls ./file
smw:/var/tmp# su - crayadm
crayadm@smw:~> cd /usr/bin
crayadm@smw:/usr/bin> ./command
```

And here are the same prompts as they appear in this publication:

```
smw# cd /etc
smw# cd /var/tmp
smw# ls ./file
smw# su - crayadm
crayadm@smw> cd /usr/bin
crayadm@smw> ./command
```

Typographic Conventions

<i>Monospace</i>	Indicates program code, reserved words, library functions, command-line prompts, screen output, file/path names, and other software constructs.
Monospaced Bold	Indicates commands that must be entered on a command line or in response to an interactive prompt.
<i>Oblique or Italics</i>	Indicates user-supplied values in commands or syntax definitions.

Proportional Bold	Indicates a GUI Window , GUI element , cascading menu (Ctrl → Alt → Delete), or key strokes (press Enter).
\ (backslash)	At the end of a command line, indicates the Linux® shell line continuation character (lines joined by a backslash are parsed as a single line).

Feedback

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

Trademarks

The following are trademarks of Cray Inc. and are registered in the United States and other countries: CRAY and design, SONEXION, URIKA, and YARCDATA. The following are trademarks of Cray Inc.: APPRENTICE2, CHAPEL, CLUSTER CONNECT, CRAYDOC, CRAYPAT, CRAYPORT, DATAWARP, ECOPHLEX, LIBSCI, NODEKARE. The following system family marks, and associated model number marks, are trademarks of Cray Inc.: CS, CX, XC, XE, XK, XMT, and XT. The registered trademark LINUX is used pursuant to a sublicense from LMI, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis. Other trademarks used in this document are the property of their respective owners.

1.1 Related Publications

This publication supersedes *SMW HA Initial Installation Guide for XC Series Systems* (SLEHA12.SP0.UP01), *SMW HA Installation Guide* (S-0044-F), and *Installing, Configuring, and Managing SMW Failover on the Cray XC System* (S-0044-D).

This publication supplements the installation procedures for a system with a single SMW in *XC™ Series Software Installation and Configuration Guide*. Only the HA-specific procedures are found in this guide.

The following publications contain additional information that may be helpful. With the exception of the first two on the list, these and other Cray publications can be found at <http://pubs.cray.com>.

- *SMW HA Release Errata* and the *SMW HA README*, which are provided with the SMW HA software release package
- *CLE Release Errata* and the *CLE README*, which are provided with the CLE release software
- *SMW Release Errata* (includes notice of any patches) and the *SMW README*, which are provided with the SMW release software
- *XC™ Series Software Installation and Configuration Guide* (S-2559)
- *XC™ Series SMW HA Administration Guide* (S-2551)
- SUSE Linux Enterprise High Availability (SLEHA) Extension 12 documentation, which provides information on the SUSE HA software, the Pacemaker Cluster Resource Manager (CRM), and related tools. SUSE manuals can be found in the `docu` directory of the SLEHA installation media, or in the directory `/usr/share/doc/` on the installed system (if installed).
- *XC™ Series System Administration Guide* (S-2393)
- *XC™ Series Configurator User Guide* (S-2560)
- *XC™ Series Lustre® Administration Guide* (S-2648)

- *XC™ Series Power Management Administration Guide (S-0043)*
- *XC™ Series System Environment Data Collections (SEDC) Guide (S-2491)*
- *XC™ Series DataWarp™ Installation and Administration Guide (S-2564), which supersedes DataWarp Installation Guide (S-2547)*
- *Cray Compiling Environment Release Overview and Installation Guide*
- *XC™ Series eLogin Installation Guide (S-2556)*
- *XC™ Series SEC Configuration Guide (S-2542) (Simple Event Correlator)*
- *XC™ Series Aries™ Network Resiliency Guide (S-0041)*
- *XC™ Series DVS Administration Guide (S-0005)*

1.2 Distribution Media

The Cray SMW SLEHA release includes one DVD or ISO file that contains the Cray SMW HA software package.

SLE HA software	<code>SLE-12-HA-DVD-x86_64-GM-CD1.iso</code>
SMW HA release	<code>smwha-sleha12sp0-12.0.3079-201701241124.iso</code>

For an initial installation and most upgrade/update installations, you will also need the release media for the operating system, SMW software, and CLE software.

For more information, see the SMW HA README file provided with the SMW HA release package. Also see the release notes and README files that are provided with the SMW and CLE release packages

2 Install and Configure an SMW HA System

This guide includes the following procedures for installing and configuring an SMW HA system:

1. [Prepare for an SMW HA Fresh Install](#) on page 9
2. [Install and Configure the XC System on the First SMW](#) on page 10
3. [Prepare to Install SMW HA Software](#) on page 10 (gather all site-customization and local changes made on the first SMW)
4. [Install the SMWHA Software on the First SMW](#) on page 12
5. [Install and Configure the Second SMW](#) on page 14
6. [Reboot the Second SMW and Power On the First SMW](#) on page 52
7. [Configure the SMW HA Cluster](#) on page 53
8. [Change Default HA Passwords After Installation](#) on page 65
9. [Configure Failover Notification](#) on page 66
10. [Configure the Power Management Database with DRBD for SMW HA](#) on page 68
11. [Finish Configuring the SMW HA System](#) on page 72

2.1 Prepare for an SMW HA Fresh Install



WARNING: When a fresh install is performed on a system, disks are wiped clean. To prevent loss of necessary data, before beginning any installation procedures, consider what configuration files, log files, or other files should be preserved, and save them in a location unaffected by the installation.

In preparation for a fresh install, do the following:

- Read the *SMW HA Release Errata* and the *SMW HA README* provided with the SMW HA release package for any additional installation-related requirements, corrections to this installation guide, and other relevant information about the release package.
- Read the Field Notices (FN) related to Cray patches to identify any required patches for this release package.
- Read the Field Notices (FN) related to kernel security fixes to identify any changes to this release package. Apply any needed changes before installing the new software.
- Verify that the network connections are in place.
- Know which configuration values are site-specific and which are defaults.
- Be familiar with the default passwords used during the installation process.

2.2 Install and Configure the XC System on the First SMW

To install and configure the XC system on the first SMW, use chapters 1 through 3 of *XC™ Series Software Installation and Configuration Guide (CLE 6.0.UP03) S-2559*. The procedure for the first SMW is the same as for a system with a stand-alone SMW, except for these differences:

- **Unmanaged interfaces.** During configuration, ensure that eth0 (interface to customer network) and eth3 (interface to admin network) are configured as unmanaged interfaces in the `cray_global_net` configuration template.
- **Power Management.** If the Cray SMWHA software will be installed immediately after installing and configuring the first SMW, skip the procedures to configure Power Management on the first SMW. Power Management for the SMW HA system will be configured later in the HA fresh install process, when the PostgreSQL Power Management Database (PMDb) will be configured as a distributed replicated block device (DRBD) connected via eth5 on both SMWs.

IMPORTANT: If the system will run with a stand-alone SMW before SMWHA is installed, configure Power Management as directed for a system with a single SMW. After installing and configuring the SMW HA system, the existing PostgreSQL PMDb data must be migrated to a temporary disk then back to the SMW disks using a special migration procedure.

- **Patch sets.** Note all patch sets that are applied on the first SMW. The second SMW must have exactly the same patch sets.
- **Customizations.** Record all site customization and local changes that are made on the first SMW. The same changes must be made on the second SMW.

When finished with the installation and configuration of the first SMW, return to this guide and continue the SMW HA installation process with [Prepare to Install SMW HA Software](#) on page 10.

2.3 Prepare to Install SMW HA Software

To prepare for installing the SMW HA system, perform the following tasks the order listed:

1. [Record All Site Customization and Local Changes to the First SMW](#) on page 10
2. [Power Off the Second SMW](#) on page 11
3. [Verify that eth0 and eth3 are Unmanaged Interfaces](#) on page 11
4. Required: Shut down the Cray system (service and compute nodes).

2.3.1 Record All Site Customization and Local Changes to the First SMW

Procedure

1. Gather information about the site customization changes that were made on the first SMW. The same changes must be made on the second SMW.

2. Record all local changes made after site customization, so that these changes can be replicated on the second SMW.

2.3.2 Power Off the Second SMW

Procedure

1. Power off the second SMW before beginning the SMW HA installation.
2. Ensure that only one SMW is powered at a time during this installation process. This is necessary to prevent device contention with shared boot RAID volumes.

2.3.3 Verify that eth0 and eth3 are Unmanaged Interfaces

About this task

For an SMW HA system, eth0 and eth3 must be configured as unmanaged interfaces in the `cray_global_net` config set worksheet so that Ansible does not manage them. Those configuration settings were set during configuration of the first SMW.

Procedure

1. Verify the configuration settings for eth0 and eth3 before installing or updating the SMW HA software.
2. Change to the global worksheet directory.

```
smw# cd /var/opt/cray/imps/config/sets/global/worksheets
```

3. Search for **customer_ethernet.unmanaged_interface** and ensure that it is set to 'true.'

```
smw# grep "customer_ethernet.unmanaged_interface" cray_global_net_worksheet.yaml
cray_global_net.settings.hosts.data.primary_smw.interfaces.customer_ethernet.unmanaged_interface: true
```

4. Search for **admin_ethernet.unmanaged_interface** and ensure that it is set to 'true.'

```
smw# grep "admin_interface.unmanaged_interface" cray_global_net_worksheet.yaml
cray_global_net.settings.hosts.data.primary_smw.interfaces.admin_interface.unmanaged_interface: true
```

5. If either of these settings is set to 'false,' change it to 'true' using the configurator interactively.
 - a. Set the value of the `unmanaged_interface` setting to 'true.'
To make the change, use the "Change a Multival Setting Field during a Configurator Session" procedure (under "Common Tasks When Using the Configurator Interactively") in *XC™ Series Configurator User Guide* (S-2560).
 - b. Inspect the contents of the following files on **both** SMWs to see if they have been modified by an Ansible play.

If any of these files have been changed by Ansible, restore them to their original contents.

```
/etc/sysconfig/network/ifcfg-eth0
/etc/sysconfig/network/ifcfg-eth1
/etc/sysconfig/network/ifcfg-eth3
```

Trouble? Contact Cray Customer Support if assistance is needed restoring these files to their original contents.

2.4 Install the SMWHA Software on the First SMW

Prerequisites

(REQUIRED) Shut down the Cray system (service and compute nodes) before using this procedure to install and configure the SMWHA software.

This procedure requires the following ISOs:

- SLE-12-HA-DVD-x86_64-GM-CD1.iso
- smwha-sleha12sp0-12.0.3079-201701241124.iso

Procedure

START A TYPESCRIPT FILE

1. Log in as root to the first SMW.
2. (First time only) Create a release directory for the typescript file.

```
smw1# mkdir -p /var/adm/cray/release
```

3. Change to the release directory.

```
smw1# cd /var/adm/cray/release
```

4. Set a variable equal to today's date.

```
smw1# export TODAY=`date +%Y%m%d`
echo $TODAY
```

5. Start a typescript file.

```
smw1# script -af ${TODAY}.suffix
```

For *suffix*, substitute a unique string to distinguish among typescript files, such as `install.1` or `update.2`.

6. Change prompt to include a timestamp.

```
smw1# PS1="\u@\h:\w \t# "
```

COPY ISOs

7. Copy the SLEHA ISO, `SLE-12-HA-DVD-x86_64-GM-CD1.iso`, to the `/root/isos` directory.
8. Copy the SMWHA release ISO, `smwha-sleha12sp0-12.0.3079-201701241124.iso`, to the `/root/isos` directory.

SAVE CURRENT SNAPSHOT NAME

9. Save the name of the current snapshot.

```
smw1# export SNAPSHOT=$(snaputil list |grep ^cur| awk '{print $2}')
```

```
smw1# echo $SNAPSHOT
```

Record this snapshot name offline (not on the SMW) so that it will be accessible later during the installation of the second SMW (smw2) while this SMW (smw1) is powered off.

SET HA SNAPSHOT NAME

10. Set variable for HA snapshot name on first SMW. Select a new target snapshot for the final HA installation target using the currently booted snapshot name.

```
smw1# export SNAPSHOT_HA=$(snaputil list |grep ^cur| awk '{print $2}')-ha
```

11. Record the snapshot name. The exact same snapshot name **must** be used on the second SMW.

IMPORTANT: Both SMWs **must** use the same snapshot name. For SMW HA systems, it is **very important** to use the exact same snapshot name for both SMWs. Failing to do this will result in HSS database (MySQL) inconsistencies between the snapshots.

```
smw1# echo $SNAPSHOT_HA
```

INSTALL HA SOFTWARE

12. Install the SMWHA software on first SMW.

```
smw1# mkdir -p /media/SMWHA
```

```
smw1# mount -o loop,ro /root/isos/smwha-sleha12sp0-12.0.3079-201701241124.iso \
```

```
/media/SMWHA
```

```
smw1# /media/SMWHA/SMWHAinstall --target $SNAPSHOT_HA
```

```
smw1# snaputil default $SNAPSHOT_HA
```

COLLECT INFORMATION FOR THE SECOND SMW

13. Record information and save the files required for the second SMW.
 - a. Record the final HA snapshot name used for the first SMW. The second SMW must use the same final snapshot name.
 - b. Copy the `cray_bootraid_config.yaml` file to a remote system.

```
smw1# scp -p \
```

```
/var/opt/cray/imps/config/sets/global/config/cray_bootraid_config.yaml user@host:~/.
```

- c. Copy the `/var/adm/cray/install.cle.conf` file to a remote system.

```
smw1# scp -p /var/adm/cray/install.cle.conf user@host:~/.
```

- d. Record the HSS data store (MariaDB) root password, if changed on the first SMW.
- e. Record other changed passwords on the first SMW.

Table 1. Default Passwords for an SMW HA System

ID	Default Password
root on smw1	initial0
root on smw2	initial0
root (iDRAC) on smw1	initial0
root (iDRAC) on smw2	initial0
hacluster (for logging in to crm_gui)	same as SMW root (set during HA configuration)
stonith-1 resource	same as iDRAC root (set during HA configuration)
stonith-2 resource	same as iDRAC root (set during HA configuration)

- f. Record any site customization and local changes that were done on the first SMW. These changes must be duplicated exactly on the second SMW.

SHUT DOWN CLE AND POWER OFF THE FIRST SMW

14. Shut down CLE if it is running.

15. Power down the first SMW.

```
smw1# shutdown -h now
```

IMPORTANT: After the first SMW has been powered off, do not turn it back on until directed to do so during the cluster configuration procedure (after all software has been installed on the second SMW).

2.5 Install and Configure the Second SMW

Prerequisites

Before installing the second SMW, ensure that the first SMW has been powered down.

About this task

Perform the following steps in this order on the second SMW:

Procedure

1. [Install the Base Operating System on the Second SMW](#) on page 15
2. [Install SMW and CLE Software on the Second SMW](#) on page 41
3. [Configure the Second SMW for CLE System Hardware](#) on page 49

4. [Install the SMWHA Software on the Second SMW](#) on page 51

2.5.1 Install the Base Operating System on the Second SMW

At this point, install the base operating system on the second SMW using the following procedures. Because these are generic procedures, the system prompt is shown as `smw#`, not `smw2#`.

2.5.1.1 Install the Base Operating System on the SMW

The base operating system must be installed on the SMW before the Cray SMW and CLE software release packages can be installed. Cray provides two rack-mount SMW models: the Dell PowerEdge™ R815 Rack Server and the Dell PowerEdge™ R630 Rack Server. Earlier desktide SMW hardware is not supported. The figure below shows an easy way to distinguish between the two rack-mount models when viewing them from the front.

Figure 1. Distinguishing Features of Dell R815 and R630 Servers



Dell R815: 2U high and 6 drive bays



Dell R630: 1U high and 8 drive bays

Continue the installation process with [Prepare to Install the Base Linux Distribution](#) on page 15.

2.5.1.1.1 Prepare to Install the Base Linux Distribution

About this task

A full initial installation begins with installing the base operating system. This procedure provides initial steps that are common to installing the base OS on both Dell R815 and R630 SMW models.

Procedure

1. Disconnect the SMW connection to the boot RAID.
Disconnect the data cables and place protective covers on the fibre optic cable connectors (if present).
2. Connect the SMW keyboard, monitor, and mouse.
Connect a keyboard, monitor, and mouse to the USB and monitor connectors on the SMW, if not already connected.

NOTE: Once the iDRAC has been configured, the keyboard, monitor, and mouse can be connected to the iDRAC for remote console activities instead of being directly connected to the SMW console.

As the next step in preparing to install the base OS, do one of the following, depending on the SMW model.

- For a Dell R630 SMW, first configure the SMW RAID, then configure the BIOS and iDRAC:
 1. [Configure the Dell R630 SMW RAID Virtual Disks](#) on page 22
 2. [R630 SMW: Change the BIOS and iDRAC Settings](#) on page 25
- For a Dell R815 SMW, just configure the BIOS and iDRAC:
 1. [R815 SMW: Change the BIOS and iDRAC Settings](#) on page 16

After the BIOS and iDRAC settings have been configured, all SMW internal disks that are not to receive the base operating system should be physically ejected from SMW internal disk drive bays.

[2.5.1.1.1.1 R815 SMW: Change the BIOS and iDRAC Settings](#)

Prerequisites

This procedure assumes that the SMW is disconnected from the boot RAID and connected to a keyboard, monitor, and mouse.

About this task

This procedure changes the system setup for a Dell R815 SMW: the network connections, remote power control, and the remote console. Depending on the server model and version of BIOS configuration utility, there could be minor differences in the steps to configure the system. For more information, refer to the documentation for the Dell server used at this site. Because Cray ships systems with most of the installation and configuration completed, some of the steps may have been done already.

For a Dell R630 SMW, see [R630 SMW: Change the BIOS and iDRAC Settings](#) on page 25.

Procedure

1. Remove SMW non-boot internal drives.

Eject all the internal disk drives from the SMW except for the primary boot disk in slot 0 and the secondary boot disk in slot 1.
2. 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 **System Setup** menu appears.

3. Change system time.

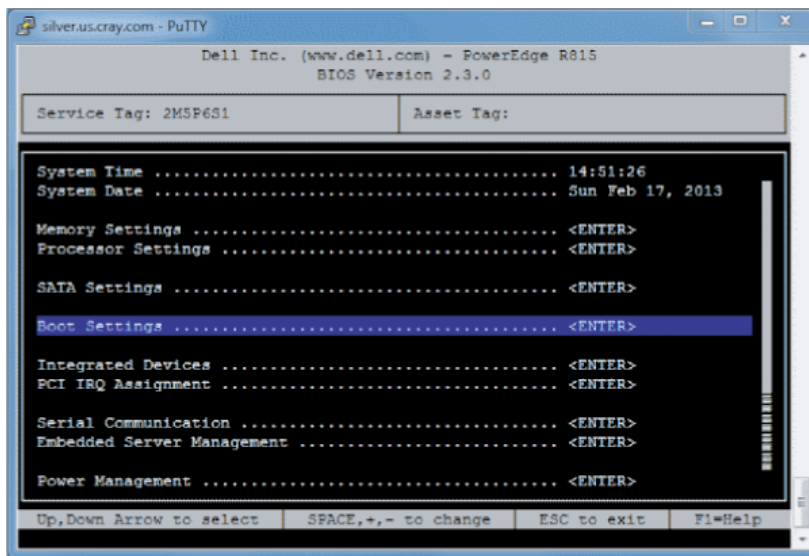
The system time should be in UTC, not in the local timezone.

 - a. Select **System Time** in the **System Setup** menu.

The hours will be highlighted in blue.

- b. Set the correct time.
 1. Press the space key to change hours.
 2. Use the right-arrow key to select minutes, then change minutes with the space key.
 3. Use the right-arrow key to select seconds, then change seconds with the space key.
 - c. Press **Esc** when the correct time is set.
4. Change boot settings.
- a. Select **Boot Settings** in the **System Setup** menu, then press **Enter**.

Figure 2. Dell R815 SMW Boot Settings Menu

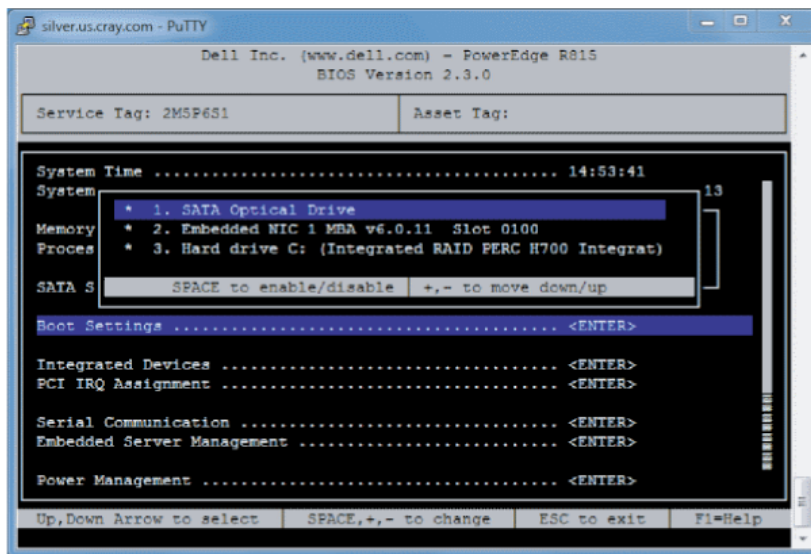


A pop-up menu with the following list appears:

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


- b. Select **Boot Sequence**, then press **Enter**.

Figure 3. Dell R815 SMW Boot Sequence Settings



- c. Change the order of items in the **Boot Sequence** list so that the optical (DVD) drive appears first, then the hard drive. If **Embedded NIC** appears in the list, it should end up below the optical drive and hard drive in the list.
 - d. Disable embedded NIC.
Select **Embedded NIC** and press **Enter**, then use the space key to disable it.
 - e. Press **Esc** to exit the **Boot Sequence** menu.
 - f. Press **Esc** again to exit the **Boot Settings** menu.
5. Change serial communication.
 - a. Select **Serial Communication** in the **System Setup** menu, then press **Enter**.
 - b. Confirm these settings in the **Serial Communication** menu.
 - **Serial Communication** is set to **On with Console Redirection via COM2**
 - **Serial Port Address** is set to **Serial Device1=COM2, Serial Device2=COM1**
 - **External Serial Connector** is set to **Serial Device2**
 - **Failsafe Baud Rate** is set to **115200**
 - c. Press **Esc** to exit the **Serial Communication** menu.
 6. Select **Embedded Server Management** in the **System Setup** menu, then press **Enter**.

Figure 4. Dell R815 SMW Embedded Server Management Settings



- a. Set **Front-Panel LCD Options** to **User-Defined LCD String** in the **Embedded Server Management** menu. Use the space key to cycle through the choices, then use the down-arrow key.
- b. Set **User-Defined LCD String** to the login hostname (e.g., cray-drac), then press **Enter**.
- c. Press **Esc** to exit the **Embedded Server Management** menu.

7. Insert base operating system DVD into SMW.

Insert the base operating system DVD labeled Cray-SMWbase12-201511021655 into the DVD drive. (The DVD drive on the front of the SMW may be hidden by a removable decorative bezel.)

8. Save BIOS changes and exit.

- a. Press **Esc** to exit the BIOS **System Setup** menu.

A menu with a list of exit options appears.

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

- b. Ensure that **Save changes and exit** is selected, then press **Enter**.

The SMW resets automatically.

9. Enter BIOS boot manager.

- a. When the BIOS POST process begins again, **quickly press the F11 key** within 5 seconds of when 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**.

10. Change the integrated Dell Remote Access Controller (iDRAC) settings.

Watch the screen carefully as text scrolls until the **iDRAC6 Configuration Utility 1.57** line is visible. When the line **Press <Ctrl-E> for Remote Access Setup within 5 sec...** displays, press **Ctrl-E** within 5 seconds.

```

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** menu appears.

11. Set iDRAC LAN to ON.**12. Configure the iDRAC LAN.**

Select **LAN Parameters**, then press **Enter**.

a. Configure iDRAC6 name.

Use the arrow key to scroll down and select **iDRAC6 Name**, then press **Enter**. Enter a value for **Current DNS iDRAC6 Name** (e.g., smw-drac), then press **Enter**.

Trouble? If unable to set the iDRAC6 name, try this:

1. Temporarily set **Register iDRAC6 Name** to "On."
2. Set **iDRAC6 Name**.
3. Return to **Register iDRAC6 Name** and set it to "Off."

b. Configure domain name.

Use the arrow key to scroll down and select **Domain Name**, then press **Enter**. Enter a value for **Current Domain Name** (e.g., us.cray.com), then press **Enter**.

c. Configure hostname string.

Use the arrow key to scroll down and select **Host Name String**, then press **Enter**. Enter a value for **Current Host Name String** (e.g., smw-drac), then press **Enter**.


d. Configure IPv4 settings.

Use the arrow key to scroll down into the **IPv4 Settings** group and confirm that the **IPv4 Address Source** is set to **static**. Then enter values for the following:

IPv4 Address (the SMW DRAC IP address)

- Subnet Mask** (the SMW iDRAC subnet mask)
- Default Gateway** (the SMW iDRAC default gateway)
- DNS Server 1** (the first site DNS server)
- DNS Server 2** (the second site DNS server)

Figure 5. Dell R815 SMW DRAC IPv4 Parameter Settings



- e. Configure IPv6 settings.

Use the arrow key to scroll down into the **IPv6 Settings** group and ensure that **IPv6** is disabled.

- f. Change the IPMI settings.

Change the IPMI settings to enable the Serial Over LAN (SOL) console.

1. Use the arrow key to scroll down to the **IPMI SETTINGS** list.
2. (SMW HA only) Ensure that **Enable IPMI over LAN** is selected. This setting is used for both SMWs in an SMW HA pair.
3. Ensure that **Channel Privilege Level Limit** is set to **Administrator**.

- g. Press **Esc** to exit **LAN Parameters** and return to the **iDRAC6 Configuration Utility** menu.

13. Configure iDRAC virtual media.

- a. Select **Domain Name**, then press **Enter**.
- b. Select **Virtual Media Configuration**, then press **Enter**.
- c. Select the **Virtual Media** line and press the space key until it indicates **Detached**.
- d. Press **Esc** to exit the **Virtual Media Configuration** menu.

14. Set the password for the iDRAC LAN root account.

Using the arrow keys, select **LAN User Configuration**, then press **Enter**. The following configuration is for both SSH and web browser access to the iDRAC.

- a. Select **Account User Name** and enter the account name "root."
- b. Select **Enter Password** and enter the intended password.

- c. Select **Confirm Password** and enter the intended password again.
- d. Press **Esc** to return to the **iDRAC6 Configuration Utility** menu.

15. Exit the iDRAC configuration utility.

- a. Press **Esc** to exit the **iDRAC6 Configuration Utility** menu.
- b. Select **Save Changes and Exit**.

The **BIOS Boot Manager** menu appears.

16. Choose to boot from SATA Optical Drive.

Using the arrow keys, select the **SATA Optical Drive** entry, then press **Enter**.

Dell R815 SMW system setup (changing default BIOS and iDRAC settings) is now complete. This completes the preparation for installing the base operating system. The next step in the process is [Install the SLES 12 Base Linux Distribution on the SMW](#) on page 36.

2.5.1.1.1.2 Configure the Dell R630 SMW RAID Virtual Disks

Prerequisites

This procedure assumes that the SMW is disconnected from the boot RAID and connected to a keyboard, monitor, and mouse.

About this task

Before installing and configuring SMW software, the base operating system needs to be installed on the SMW. And before the base operating system can be installed, the internal disk drives of the SMW must be configured as RAID virtual disks, as described in this procedure, and the default system setup for the R630 SMW node must be configured, as described in [R630 SMW: Change the BIOS and iDRAC Settings](#) on page 25.

A Dell R630 SMW has five physical disks. The SMW node must be reconfigured so that the internal Dell PERC RAID controller treats four of these disks as RAID 5 with a hot spare and the fifth disk as non-RAID. This procedure describes how to do that. Because Cray ships systems with most of the installation and configuration completed, some of the steps may be needed only if changes are made to the configuration.

This procedure includes detailed steps for the DELL R630 server using the PERC H330 Mini BIOS Configuration Utility 4.03-0010. Depending on the server model and version of RAID configuration utility, there could be minor differences in the steps to configure the system. For more information, refer to the documentation for the DELL PERC controller or server RAID controller software.

Procedure

1. Connect a keyboard, monitor, and mouse to the front panel USB and monitor connectors on the SMW, if not already connected.
2. Ensure all SMW internal disk drives are inserted into the SMW drive slots.
3. Power up the SMW. As the SMW node reboots, watch for the Power Edge Expandable RAID Controller section and be ready to press **Ctrl-R** when prompted.

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

TIP: In the RAID configuration utility:

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

4. Clear existing/default disk configuration, if necessary.

If any disk groups are currently defined:

- a. Select **Disk Group 0**, then press **F2**.
- b. Select **Delete Disk Group**, then press **Enter**.
- c. Select **Yes** in the pop-up confirmation window to confirm the changes.

5. Switch disk controller from HBA-Mode to RAID-Mode.

Some SMW hardware might be configured for HBA-Mode. If it is, then change it to RAID-Mode using the following substeps. If it is not, then skip these substeps.

- a. Switch disk controller from HBA-Mode to RAID-Mode.
 1. Press **Ctrl-N** (multiple times) to move to the **Ctrl Mgmt** tab.
 2. Press **Tab** (multiple times) to get to **Personality Mode**.
 3. Press **Enter** to see choice between **RAID-Mode** and **HBA-Mode**.
 4. Use the up-arrow or down-arrow key to select **RAID-Mode**, then press **Enter**.
 5. Press **Tab** (multiple times) to get to **Apply**, then press **Enter**. This message appears: "The operation has been performed successfully. Reboot the system for the change to take effect."
 6. Press **Enter**.
- b. Exit RAID configuration utility.
 1. Press **Esc** to exit the RAID configuration utility.
 2. Select **OK** to confirm, then press **Enter**.
- c. Reboot the SMW.

Press **Ctrl-Alt-Delete** at the prompt to reboot. The server will restart the boot process. Be prepared to press **Ctrl-R** when prompted.
- d. Enter RAID configuration utility.

As the SMW node reboots, enter the RAID controller configuration utility by pressing **Ctrl-R** when prompted. This will return to the point prior to switching from HBA-Mode to RAID-Mode.

6. Configure most internal disks to appear as `/dev/sda` in a RAID-5 virtual disk.

This step configures most internal disks to appear as `/dev/sda` in a RAID-5 virtual disk. The R630 has four identical 500-GB drives and one 1-TB drive. The 1-TB drive will be excluded from this RAID-5 configuration. Instead, that drive will be used to hold the postgresql database with Power Management data.

- a. Select **No Configuration Present**, then press the **F2** key.

- b. Select **Convert to RAID capable**, then press **Enter**. The **Convert Non-RAID Disks to RAID capable** screen appears.
 - c. Create virtual disk /sda
7. Convert non-RAID disks to RAID-capable.
 - a. Press **Enter** to check the box for a physical disk, which selects it for this RAID-5 disk group. This action also advances the selection to the next disk. In this manner, select all four of the identical 500-GB drives but exclude the 1-TB drive (leave it unselected).
 - b. Press **Tab** to move to **OK**, then press **Enter**.
8. Verify the virtual disk changes.

To verify the virtual disk changes, compare settings with those shown in the figure.
9. Create virtual disk sda.
 - a. Use up-arrow to return to the **No Configuration Present!** item.
 - b. Press **F2** to see a pop-up menu.
 - c. Press **Enter** to choose **Create New VD**.

The **Convert Non - RAID Disks to RAID capable** screen appears. The only disk left on this screen should be the large (1-TB) disk which was excluded earlier. It should not be added to the RAID capable set of disks, so continue to exclude it.
 - d. Press **Tab** to move from the list of disks to **Cancel**, then press **Enter**.

This cancels the conversion of non-RAID disks to RAID capable. The **Create New VD** screen appears.
10. Create new virtual disk (VD).
 - a. Press **Enter** to switch from **RAID-0** to other options.
 - b. Use down-arrow to select **RAID-5**, then press **Enter**.
 - c. Press **Tab** to move to the **Physical Disks** area.
 - d. Press **Enter** to select each disk except one.

One disk should not be selected so that it can become the hot spare (configured in the next step).
 - e. Press **Tab** to move to **VD Name**.
 - f. Select name sda.
 - g. Press **Tab** to move to **Advanced**, then press **Enter**.

The **Create Virtual Disk-Advanced** screen appears.
11. Configure one disk as the hot spare.
 - a. Press **Tab** multiple times to move to **Initialize**, then press **Enter** to select it.

A pop-up window with the following text appears: "Initialization will destroy data on the virtual disk. Are you sure you want to continue?"
 - b. Press **Tab** or arrow keys to move to **OK**, then press **Enter** to confirm initialization.

- c. Press **Tab** to move to **Configure HotSpare**, then press **Enter** to select it.
- d. Press **Tab** or arrow keys to move to **OK** on the **Create Virtual Disk-Advanced** screen, then press **Enter**.
- e. Press **Tab** or arrow keys to move to **OK** on the **Create New VD** screen, then press **Enter**.
A pop-up window with the following text appears: "Virtual disk is successfully created and initialized."
- f. Press **Tab** or arrow keys to move to **OK**, then press **Enter**.
A pop-up window with the following text appears: "Dedicated Hotspare for Disk Group 0."
- g. Select the disk to be the hot spare, then press **Enter**.
- h. Press **Tab** or arrow keys to move to **OK**, then press **Enter**.
A pop-up window with the following text appears: "Initialization complete on VD 0."
- i. Press **Tab** or arrow keys to move to **OK**, then press **Enter**.

The RAID will initialize in the background.

12. Exit RAID configuration utility.

Exit the RAID configuration utility, reboot, and then begin installing the base operating system.

- a. Press the **Esc** key to exit the RAID configuration utility.
- b. Select **OK**, then press **Enter** to confirm.

13. Reboot the system.

A message appears that prompts to reboot.

ATTENTION: Only the disk drives configured to be the RAID-5 virtual disk sda should be inserted into the SMW internal drive bays when installing SLES 12.

- a. Eject the 1-TB disk (which was not added to the RAID-5 virtual disk sda) from the SMW.

This will be re-inserted when SLES 12 installation is complete.

- b. Press **Ctrl-Alt-Delete**.

The server will restart the boot process and will not interrupt RAID initialization. During the system reboot, be prepared to press **F2** when prompted, to change the system setup.

RAID configuration is now complete. The next step in preparing to install the base operating system is to configure the system setup for the R630 SMW node, as described in [R630 SMW: Change the BIOS and iDRAC Settings](#) on page 25.

[2.5.1.1.1.3 R630 SMW: Change the BIOS and iDRAC Settings](#)

Prerequisites

This procedure assumes that the internal disk drives of the SMW have just been configured as RAID virtual disks and the system is rebooting. If the system is not rebooting, press **Ctrl-Alt-Delete** to reboot.

About this task

This procedure describes how to change the system setup for the SMW: the network connections, remote power control, and the remote console. This procedure includes detailed steps for the Dell R630 server. Depending on the server model and version of BIOS configuration utility, there could be minor differences in the steps to

configure the system. For more information, refer to the documentation for the Dell server used at this site. Because Cray ships systems with most of the installation and configuration completed, some of the steps may have been done already.

For a Dell R815 server, see [R815 SMW: Change the BIOS and iDRAC Settings](#) on page 16.

Procedure

Watch as the system reboots and the BIOS power-on self-test (POST) process begins. Be prepared to press **F2**, when prompted, to change the system setup.

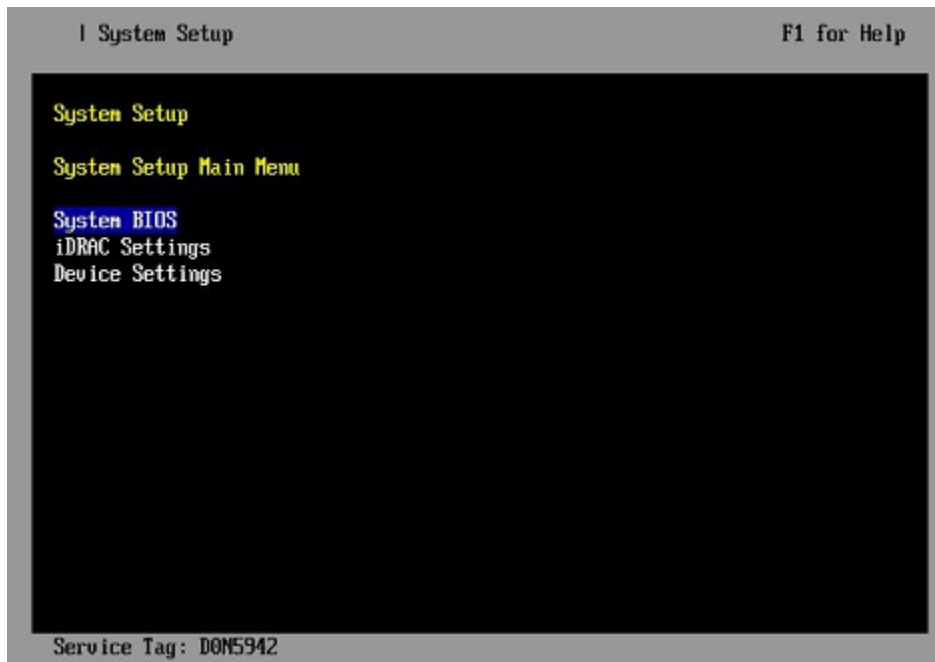
1. Press the **F2** key immediately after the following messages appear in the upper-left of the screen:

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

When the **F2** keypress is recognized, the **F2 = System Setup** line changes color from white-on-black to white-on-blue.

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

Figure 6. Dell R630 System Setup Main Menu



TIP: In system setup screens,

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

2. Change the BIOS settings.

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

The **System BIOS Settings** screen appears.

Figure 7. Dell R630 System BIOS Settings Screen



b. Change Boot Settings.


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

Figure 8. Dell R630 Boot Settings Screen




2. Ensure that **Boot Mode** is **BIOS** and not **UEFI**.
3. Select **Boot Option Settings**, then press **Enter**.
4. Select **Boot Sequence** on the **Boot Option Settings** screen, then press **Enter** to view a pop-up window with the boot sequence.

Figure 9. Dell R630 BIOS Boot Sequence



5. Change the boot order in the pop-up window so that the optical drive appears first, then the hard drive. If **Integrated NIC** appears in the list, it should end up below the optical drive and hard drive in the list.
TIP: Use the up-arrow or down-arrow key to highlight or select an item, then use the **+** and **-** keys to move the item up or down.
 6. Select **OK**, then press **Enter** to accept the change.
 7. Click the box next to **Hard drive C:** under the **Boot Option/Enable/Disable** section to enable it. Do the same for the optical drive, if necessary.
 8. Select **integrated NIC**, then press **Enter** to disable it.
 9. Press **Esc** to exit **Boot Option Settings**.
 10. Press **Esc** to exit **Boot Settings** and return to the **System BIOS Settings** screen.
- c. Change Serial Communication Settings.

Figure 10. Dell R630 System BIOS Settings: Serial Communication




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

Figure 11. Dell R630 Serial Communication Screen



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

Figure 12. Dell R630 iDRAC Settings Screen



4. Change the iDRAC network.

- a. Select **Network** to display a long list of network settings.
- b. Change the DNS DRAC name.

Use the arrow key to scroll down to **DNS DRAC Name**, then enter an iDRAC hostname that is similar to the SMW node hostname (e.g., cray-drac).

- c. Change the static DNS domain name.

Use the arrow key to scroll down to **Static DNS Domain Name**, then enter the DNS domain name and press **Enter**.

- d. Change the IPv4 settings.

Use the arrow key to scroll down to the **IPV4 SETTINGS** list.

1. Ensure that IPv4 is enabled.

- a. If necessary, select **Enable IPV4**, then press **Enter**.

- b. Select **<Enabled>** in the pop-up window, then press **Enter** to return to the previous screen.

2. Ensure that DHCP is disabled.

- a. If necessary, select **Enable DHCP**, then press **Enter**.

- b. Select **<Disabled>** in the pop-up window, then press **Enter** to return to the previous screen.

3. Change the IP address.

- a. Select **Static IP Address**.

- b. Enter the IP address of the iDRAC interface (`ipmi0`) for the SMW, then press **Enter**.

4. Change the gateway.

- a. Select **Static Gateway**.
 - b. Enter the appropriate value for the gateway of the network to which the iDRAC is connected, then press **Enter**.
5. Change the subnet mask.
 - a. Select **Subnet Mask**.
 - b. Enter the subnet mask for the network to which the iDRAC is connected (such as 255.255.255.0), then press **Enter**.
6. Change the DNS server settings.
 - a. Select **Static Preferred DNS Server**, enter the IP address of the primary DNS server, then press **Enter**.
 - b. Select **Alternate DNS Server**, enter the IP address of the alternate DNS server, then press **Enter**.
- e. Change the IPMI settings.

Change the IPMI settings to enable the Serial Over LAN (SOL) console.

 1. Use the arrow key to scroll down to the **IPMI SETTINGS** list.
 2. Ensure that **Enable IPMI over LAN** is selected.

TIP: Use the left-arrow or right-arrow to switch between two settings.
 3. Ensure that **Channel Privilege Level Limit** is set to **Administrator**.
- f. Exit Network screen.

Press the **Esc** key to exit the **Network** screen and return to the **iDRAC Settings** screen.
5. Change hostname in iDRAC LCD display.

Change front panel security to show the hostname in LCD display.

 - a. Use the arrow key to scroll down and highlight **Front Panel Security** on the **iDRAC Settings** screen, then press **Enter**.
 - b. Select **Set LCD message**, then press **Enter**.
 - c. Select **User-Defined String**, then press **Enter**.
 - d. Select **User-Defined String**, then enter the SMW hostname and press **Enter**.
 - e. Press the **Esc** key to exit the **Front Panel Security** screen.
6. (Optional) Change the iDRAC **System Location** fields.

Change the **System Location** configuration on the **iDRAC Settings** screen to set any of these fields: **Data Center Name**, **Aisle Name**, **Rack Name**, and **Rack Slot**.
7. Configure iDRAC virtual media.
 - a. Select **Domain Name**, then press **Enter**.
 - b. Select **Virtual Media Configuration**, then press **Enter**.
 - c. Select the **Virtual Media** line and press the space key until it indicates **Detached**.
 - d. Press **Esc** to exit the **Virtual Media Configuration** menu.

8. Set the password for the iDRAC root account.
 - a. Use the arrow key to highlight **User Configuration** on the **iDRAC Settings** screen, then press **Enter**.
 - b. Confirm that User Name is root. Select **User Name**, then enter the "root" user name.
 - c. Select **Change Password**, then enter a new password.
 - d. Reenter the new password in the next pop-up window to confirm it (the default password is "calvin").
 - e. Press the **Esc** key to exit the **User Configuration** screen.

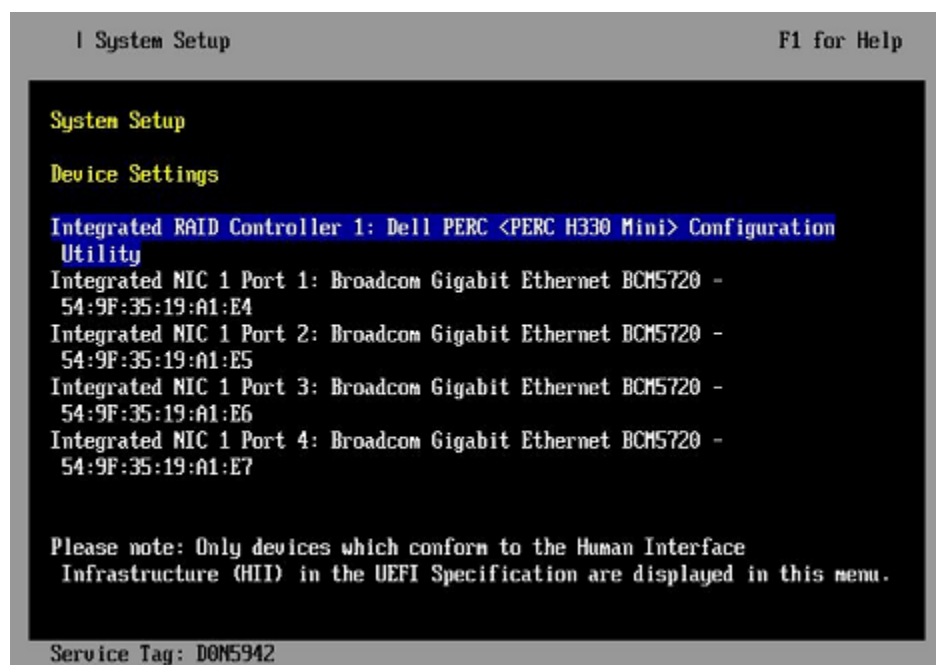
9. Exit iDRAC settings.
 - a. Press the **Esc** key to exit the **iDRAC Settings** screen.
A "Settings have changed" message appears.
 - b. Select **Yes**, then press **Enter** to save the changes.
A "Success" message appears.
 - c. Select **Ok**, then press **Enter**.
The main screen (**System Setup Main Menu**) appears.

10. Change device settings.

These steps disable an integrated NIC device by changing the setting for the integrated NIC on a port from **PXE** to **None**.

- a. Change Integrated NIC 1 Port 1
 1. Select **Device Settings** on the **System Setup Main Menu**, then press **Enter**. The **Device Settings** screen appears.

Figure 13. Dell R630 Device Settings Screen




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

Figure 14. Dell R630 BIOS MBA Configuration Settings



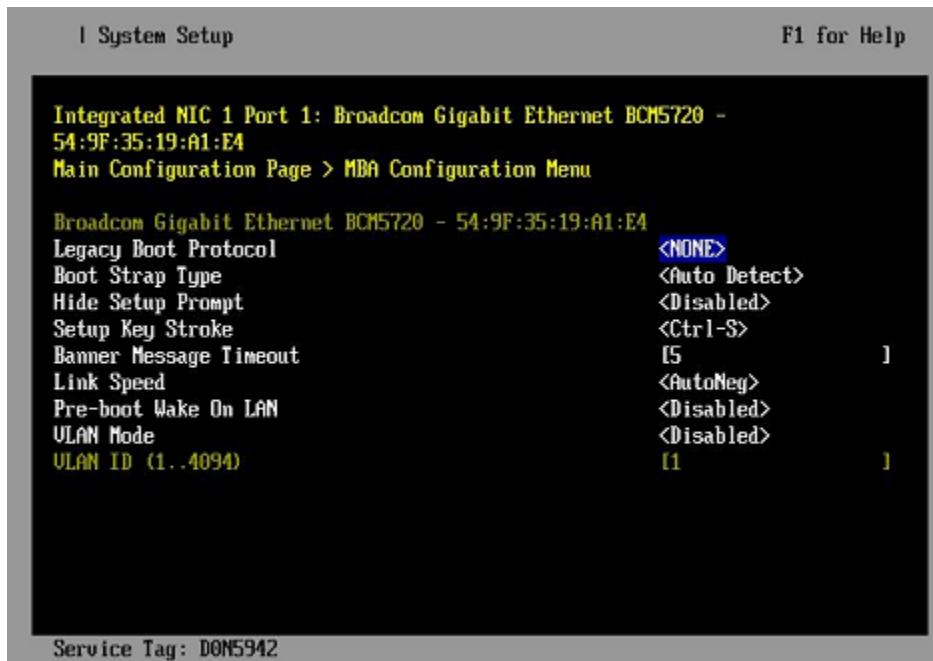
4. Select **Legacy Boot Protocol** on the **MBA Configuration Menu** screen, use the right-arrow or left-arrow key to highlight **None**, then press **Enter**.
 5. Press the **Esc** key to exit the **MBA Configuration Menu** screen.
 6. Press **Esc** to exit the **Main Configuration Page** screen. A "Warning Saving Changes" message appears.
 7. Select **Yes**, then press **Enter** to save the changes. A "Success" message appears.
 8. Select **OK**, then press **Enter**. The **Device Settings** screen appears.
 9. Press **Esc** to exit the **Device Settings** screen. A "Settings have changed" message appears.
 10. Select **Yes**, then press **Enter** to save the changes. A "Settings saved successfully" message appears.
 11. Select **OK**, then press **Enter**. The main screen (**System Setup Main Menu**) appears.
- b. Change Integrated NIC 1 Port 2
1. Select **Device Settings** on the **System Setup Main Menu**, then press **Enter**. The **Device Settings** screen appears.

Figure 15. Dell R630 Device Settings Screen



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

Figure 16. Dell R630 BIOS MBA Configuration Settings



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

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

SMW system setup (changing default BIOS and iDRAC settings) is now complete. This completes the preparation for installing the base operating system. The next step in the process is [Install the SLES 12 Base Linux Distribution on the SMW](#) on page 36.

2.5.1.1.2 Install the SLES 12 Base Linux Distribution on the SMW

Prerequisites

This procedure assumes the following:

- The BIOS and iDRAC settings have just been changed on the SMW and it is restarting the boot process.
- All SMW internal disks that are not to receive the operating system are physically ejected from SMW internal disk drive bays.
- All connections to the boot RAID are unplugged so that no disk devices from the boot RAID will inadvertently lose existing data or receive the operating system.

About this task

This procedure describes the base operating system installation process. It provides detailed instructions for installing SLES 12 on the SMW (both Dell R815 and R630 models); configuring the SMW; and performing final steps: reconnect cables, reinsert drives, and reboot the SMW. To install the base operating system, use the DVD labeled Cray-SMWbase12-201511021655, which contains SUSE Linux Enterprise Server version 12 (SLES 12).

Procedure

SLES 12 SOFTWARE PACKAGE INSTALLATION

1. Select one of the **Cray SMW Initial Install** options.

Within 10 to 15 seconds after this **SUSE Linux Enterprise Server** boot menu displays, use the arrow key to scroll down and select one of the install options, then press **Enter**.

```
- Boot from Hard Disk
- Cray SMW Initial Install with software RAID1
- Cray SMW Initial Install without software RAID
- Rescue System
- Check Installation Media
- Firmware Test
- Memory Test
```

Select the option that is best for the SMW model:

For a Dell R815 SMW

Select **Cray SMW Initial Install with software RAID1**, a mirrored boot disk option, which creates a software RAID1 mirror on the first two drives. This option is best for a Dell R815 because the R815 should use two disk drives to become the software RAID1 mirror.

For a Dell R630 SMW

Select **Cray SMW Initial Install without software RAID**, a non-mirrored boot disk option, for servers with a single disk or virtual disk. This option is best for a Dell R630 because the R630 should have the internal RAID controller configured to present four disk drives as a virtual disk.



WARNING: If the selection is not made in time, the system will boot from the default selection, which is **Boot from Hard Disk**. If that happens, shut down the SMW, then start the power-up sequence again.

Note: The upper left corner of the installation screen has a date/time stamp for when the bootable SLES 12 DVD was created.

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

```
Starting ... Loading Linux kernel
Initializing
Preparing System for Automated Installation
Initializing the Installation Environment
System Probing
Installation Settings
```

2. Review installation settings while the installation pauses on the **Installation Settings** screen.

3. Confirm the language for the SMW.

English (US) is the primary language by default. To change the primary language:

a. Select the **Language** heading in the **Installation Settings** screen.

The **Languages** window opens.

b. Select a language (or multiple languages) from the drop-down menu, then select **Accept** at the bottom of the window.

4. Begin automated install.

a. On the **Installation Settings** screen, select **Install**.

The **Confirm Installation** pop-up window appears.

b. Select **Install**.

The installation of software packages runs for approximately 20 minutes. The process automatically reboots the SMW from the hard disk, and the installation process continues with system configuration.

SYSTEM CONFIGURATION

5. Log in to SMW as root.

When the login screen is displayed with the `crayadm` account as the account which will be logged in:

a. Select **Not listed?**, then enter `root` for the username.

b. Either press **Enter** or select **Sign In**.

c. Enter the password for root.

6. Change default passwords on the SMW by executing the following commands.

The SMW contains its own `/etc/passwd` file that is separate from the password file for the rest of the CLE system.

```
smw# passwd root
```

```
smw# passwd crayadm
```

```
smw# passwd mysql
```

7. Change the SMW local time zone, if needed.

The default time zone on the SMW is **America/Chicago**. To change it:

a. Execute this command:

```
smw# yast2 timezone
```

`yast2` opens a new window for changing the time zone, then a pop-up window appears with this message: "file `/etc/ntp.conf` has been changed manually. YaST might lose some of the changes."

b. Select the **Do not show this message anymore** checkbox, then select **Continue**.

c. Choose the time zone either by selecting a region on the map or by using the drop-down menus for **Region** and **Time Zone**.

d. Select **Other Settings** if the time is incorrect, then select the **Manually** radio button and enter **Current Time** and **Current Date**. Select **Accept** when done.

e. Select **OK** when done with time zone settings.

8. Configure the SMW firewall.

The SUSE firewall settings may need to be adjusted to match site firewall policy and to customize for site IP addresses. These steps enable and configure the firewall.

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

a. Save the SUSE firewall configuration.

Before modifying the SUSE firewall settings, make a copy of the configuration file.

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

b. Check current firewall settings.

Check current firewall settings and change to support any site requirements. During the process of configuring Cray SMW and CLE software, some of the firewall settings may be adjusted. SSH access is one of the protocols permitted through the firewall from the external network to the SMW.

```
smw# iptables -L
```

```
smw# vi /etc/sysconfig/SuSEfirewall12
```

c. Start the firewall immediately.

Invoke the modified configuration.

```
smw# systemctl start SuSEfirewall12_init.service
smw# systemctl start SuSEfirewall12.service
```

- d. Ensure that the firewall will start at next boot.

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

```
smw# systemctl enable SuSEfirewall12_init.service
smw# systemctl enable SuSEfirewall12.service
```

- e. Verify firewall changes.

Verify the changes to the `iptables`.

```
smw# iptables -nvL
```

9. Configure LAN on the SMW.

Set network configuration for `eth0` and the hostname for the SMW.

- a. Execute this command:

```
smw# yast2 lan
```

The **Network Settings** screen appears with the **Overview** tab highlighted.

- b. Select the **eth0** line on the **Overview** tab, then select **Edit**.

The **Network Card Setup** screen appears with the **Address** tab highlighted.

- c. Select **Statically Assigned IP address** on the **Address** tab and enter values for IP address, subnet mask, and hostname (including the domain name). Then select **Next**.

- d. Select the **Hostname/DNS** tab on the **Network Settings** screen.

1. For the **Hostname and Domain Name** area, enter Hostname and Domain Name.
2. For the **Name Servers and Domain Search List**, enter Name Server 1, Name Server 2, Name Server 3, and Domain Search.

- e. Select the **Routing** tab on the **Network Settings** screen, then enter the Default IPv4 Gateway (for the network connected to `eth0`) and set Device to `eth0` using the dropdown menu.

- f. Click **OK** after all of the **Network Settings** have been prepared.

FINAL STEPS

10. Reconnect boot RAID disk cables.

Remove the protective covers from the Fibre Channel or SAS cable connectors, clean the ends of the cable connectors, and reconnect the data cables that connect the SMW to the boot RAID.

11. Reinsert SMW non-boot internal drives.

Reinsert all of the SMW internal disk drives that were removed earlier.

TIP: It is not necessary to turn off the power for the SMW before inserting these drives—the operating system can be in a booted state.

12. Eject the Cray-SMWbase12 DVD.

If the base operating system DVD (Cray-SMWbase12-201511021655) is still in the DVD drive, eject it.

```
smw# eject
```

13. Reboot the SMW.

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

```
smw# reboot
```

If the SMW was configured with RAID1, then it may still be synchronizing the data between the two disks in the RAID1 mirror. The resync can take about 30 minutes when SLES 12 is freshly installed. If the SMW is rebooted at this point in the process, that resync will be interrupted. However, that is not a problem because as soon as the SMW is up again, the resync process will continue.

(R815 SMW only) To check the status of any RAID1 resync activities on an R815 SMW, look at `/proc/mdstat`.

In this example, the resync of md127 finishes in 24.3 minutes.

```
smw# cat /proc/mdstat
Personalities : [raid1]
md125 : active raid1 sdc2[1] sda2[0]
      33559424 blocks super 1.0 [2/2] [UU]
      bitmap: 0/1 pages [0KB], 65536KB chunk

md126 : active raid1 sda1[0] sdc1[1]
      4200384 blocks super 1.0 [2/2] [UU]
      bitmap: 0/1 pages [0KB], 65536KB chunk

md127 : active raid1 sda3[0] sdc3[1]
      206437248 blocks super 1.0 [2/2] [UU]
      [=====>.....] resync = 33.7% (69700352/206437248)
      finish=24.3min speed=93748K/sec
      bitmap: 2/2 pages [8KB], 65536KB chunk

unused devices: <none>
```

- For a stand-alone SMW or the first SMW in an SMW HA system, the next step in the installation and configuration of the SMW base operating system is [Configure Boot RAID Devices](#).
- (SMW HA only) For the second SMW in an SMW HA system, there is no need to configure the boot RAID because it is shared with the first SMW and has already been configured. The next step in the process is [Make a Snapshot Manually](#) on page 40.

2.5.1.2 Make a Snapshot Manually

Prerequisites

This procedure assumes that the SLES 12 base operating system has been installed on the SMW and boot RAID devices have been configured, but no other software has been installed yet.

About this task

A Btrfs snapshot of the SMW should be created immediately after SLES 12 has been installed and before any files or directories have been modified by Cray's installation software or the rest of the installation process. With this snapshot, it will be possible to revert to this point if an initial/fresh install is repeated.

Snapshots are usually made using the `snaputil` program, but that program has not been installed at this point in the installation process. `snaputil` will be installed to the SMW with other Cray RPMs for the SMW and will be used for all Btrfs snapshot manipulations after this point.

Note that snapshots capture content in these three directories: `/`, `/var/lib/mysql`, and `/var/opt/cray/repos`. Used in conjunction with backups of config sets, they provide enough information to be able to re-create the state of the system at a particular time.

Procedure

1. Determine the root subvolume.

It will be the string starting with "UUID." In this example it is "UUID=ffb0b613-2033-4835-87b5-6ca8ff1bacde."

```
smw# grep " / " /etc/fstab
UUID=ffb0b613-2033-4835-87b5-6ca8ff1bacde /          btrfs
defaults          0 0
```

2. Mount the root subvolume.

Substitute the correct subvolume string for the example string shown in this command.

```
smw# mount -o subvol=@ UUID=ffb0b613-2033-4835-87b5-6ca8ff1bacde /mnt
```

3. Create a subvolume for snapshots (if `/mnt/snapshots` does not already exist).

```
smw# btrfs sub create /mnt/snapshots
```

4. Create the snapshot (if `/mnt/snapshots/SLES12` does not already exist).

```
smw# btrfs sub snap / /mnt/snapshots/SLES12
```

5. Unmount the snapshot.

```
smw# umount /mnt
```

6. Make a new `/media/root-sv` directory.

```
smw# mkdir -p /media/root-sv
```

7. Mount root subvolume under `/media/root-sv` instead of `/mnt` as was used above.

Substitute the correct subvolume string for the example string shown in this command.

```
smw# mount -o subvol=@ UUID=ffb0b613-2033-4835-87b5-6ca8ff1bacde /media/root-sv
```

A "SLES12" snapshot has been made. Reboot to this snapshot whenever it is necessary to restart a fresh software installation from this point.

2.5.2 Install SMW and CLE Software on the Second SMW

To install the SMW and CLE software on the second SMW, use a subset of the full procedures used on the first SMW. Follow these procedures in the order listed.

1. [Start a Typescript File on the Second/Passive SMW](#) on page 42
2. [Prepare to Bootstrap the SMW Installation on the Second SMW](#) on page 42
3. [Bootstrap the SMW Installation on the Second SMW](#) on page 45
4. [Provision SMW Storage on the Second SMW](#) on page 46
5. [Run the Installer for an Initial Installation](#) on page 46
6. [Set Default Snapshot and Boot the SMW](#) on page 48

2.5.2.1 Start a Typescript File on the Second/Passive SMW

About this task

Sites can make as few or as many typescripts as they deem useful. Cray recommends starting a typescript file on each SMW in an SMW HA system at these software installation/configuration milestones:

- just before installing a new software release
- just before configuring the newly installed software

All of these steps apply to the second or 'passive' SMW.

Procedure

1. Log in as root to the SMW.
2. (First time only) Create a release directory for the typescript file.

```
smw2# mkdir -p /var/adm/cray/release
```

3. Change to the release directory.

```
smw2# cd /var/adm/cray/release
```

4. Set a variable equal to today's date.

```
smw2# export TODAY=`date +%Y%m%d`  
smw2# echo $TODAY
```

5. Start a typescript file.

```
smw2# script -af ${TODAY}.suffix
```

For *suffix*, substitute a unique string to distinguish among typescript files, such as `install.1` or `update.2`.

6. Change prompt to include a timestamp.

```
smw2# PS1="\u@\h:\w \t# "
```

2.5.2.2 Prepare to Bootstrap the SMW Installation on the Second SMW

Prerequisites

This procedure assumes that the base operating system has been installed on the SMW and that the boot RAID has been set up.

About this task

The installer requires several ISO files to be available for setting up and installing packages from SLE repositories. The names of these ISOs are hard-coded in the installer configuration, but the containing directory can be anywhere that makes sense.

IMPORTANT: The default location for these ISO files is `/root/isos`. The `--iso-dir` argument must be specified for `SMWinstall` if this is not the correct location for the ISO files on this system.

The Cray release distribution media consists of one DVD and several other pieces of media that may be on DVDs or furnished as ISO files. These ISO files are available for download at CrayPort (<https://crayport.cray.com>).

Procedure

COLLECT ISO FILES

1. Make a directory on the SMW to hold the ISO files, if one does not already exist.

Instead of placing the ISOs directly in `/root/isos`, use these two commands to place that directory into the `btrfs` subvolume `/var/adm/cray`, which is exempt from snapshots. This prevents the large ISO files from unnecessarily increasing the size of snapshots.

```
smw2# mkdir -p /var/adm/cray/release/isos
smw2# ln -s /var/adm/cray/release/isos /root/isos
```

2. Download the SLES 12 distribution ISOs to the new directory on the SMW.
 - `SLE-12-Module-Legacy-DVD-x86_64-GM-DVD1.iso`
 - `SLE-12-Module-Public-Cloud-DVD-x86_64-GM-DVD1.iso`
 - `SLE-12-SDK-DVD-x86_64-GM-DVD1.iso`
 - `SLE-12-Server-DVD-x86_64-GM-DVD1.iso`
 - `SLE-12-WE-DVD-x86_64-GM-DVD1.iso`
 - `SLE-12-HA-DVD-x86_64-GM-CD1.iso`
3. Download the CentOS 6.5 distribution ISO (`CentOS-6.5-x86_64-bin-DVD1.iso`) to the new directory on the SMW.
4. Download CLE 6.0 and SMW 8.0 SLES12 ISOs to the new directory on the SMW.
 - SMW release: `smw-8.0.3075-201701182038.iso`
 - CLE release: `cle-6.0.3074-201701182038.iso`
5. Download the SLES12 security updates ISO (`sleupdate-sle12sp0-201611021158.iso`) to the new directory on the SMW.

6. Make a directory on the SMW to hold any patches that may be available on CrayPort, if it does not already exist.

```
smw2# mkdir -p /var/adm/cray/release/patchsets
```

7. Download SMW and CLE patches to the patchset directory on the SMW, as described in the release notes.

MOUNT MEDIA

8. Set an environment variable for and mount the SMW media.

- a. Confirm that this is the right SMW media.

```
smw2# ls -l /root/isos/smw*iso
-rw-r--r-- 1 root root 427184128 Nov  9 10:41 smw-8.0.3075-201701182038.iso
```

- b. Set environment variables for the SMW media.

Use the release string and the build date-time stamp as the values for SMW_RELEASE and SMW_SOFTWARE, as shown in this example.

```
smw2# export SMW_RELEASE=8.0.3075
smw2# echo $SMW_RELEASE

smw2# export SMW_SOFTWARE=201701182038
smw2# echo $SMW_SOFTWARE
```

- c. Mount the SMW release media.

```
smw2# mkdir -p /media/SMW
smw2# mount -o loop,ro /root/isos/smw-${SMW_RELEASE}-${SMW_SOFTWARE}.iso \
/media/SMW
```

9. Set an environment variable for the CLE media.

- a. Confirm that this is the right CLE media.

```
smw2# ls -l /root/isos/cle*iso
-rw-r--r-- 1 root root 1146388480 Nov  9 09:22 cle-6.0.3074-201701182038.iso
```

- b. Set environment variables for the CLE media.

Use the release string and the build date-time stamp as the values for CLE_RELEASE and CLE_SOFTWARE, as shown in this example.

```
smw2# export CLE_RELEASE=6.0.3074
smw2# echo $CLE_RELEASE

smw2# export CLE_SOFTWARE=201701182038
smw2# echo $CLE_SOFTWARE
```

10. Set an environment variable for the SLES12 security updates media.

Use the entire name of the SLES12 security updates media as the environment variable. This will be used when installing SMW and CLE software and SLES12 security updates together later in the process.

```
smw2# export SLE_SOFTWARE=sleupdate-sle12sp0-201611021158
smw2# echo $SLE_SOFTWARE
```

PREPARE THE INSTALL CONFIGURATION FILE

These steps use the configuration file `/var/adm/cray/install.cle.conf` that was saved from the first SMW in [Install the SMWHA Software on the First SMW](#) on page 12.

IMPORTANT: Do not turn on the first SMW at this point. If that file was not saved before turning off the first SMW, skip these two steps. Instead, wait until after the cluster has been configured and both SMWs are up, and then complete this task.

11. Retrieve the `install.cle.conf` from the first SMW.

The `install.cle.conf` file contains configuration that controls the installer's image building behavior.

```
smw2# scp -p user@host:~/install.cle.conf /var/adm/cray/install.cle.conf
```

12. Ensure that image building is disabled.

Images were built as part of the `smw1` installation process, so if the output of this command shows that `build_images` is set to `yes`, edit `/var/adm/cray/install.cle.conf` and set `build_images` to `no`.

```
smw2# grep "build_images" /var/adm/cray/install.cle.conf
build_images: no
```

2.5.2.3 Bootstrap the SMW Installation on the Second SMW

Prerequisites

This procedure assumes that the `cray_bootraid_config.yaml` file on the first SMW (`smw1`) was saved in [Install the SMWHA Software on the First SMW](#) on page 12.

About this task

This procedure runs `SMWinstall` in bootstrap mode, which installs IMPS and Ansible on the SMW, along with some of the global configuration templates.

Procedure

1. If multipath will be used, enable and start the multipath daemon now.

```
smw2# systemctl enable multipathd
```

```
smw2# systemctl start multipathd
```

2. Copy the storage configuration template from first SMW to the second SMW. This template was saved when configuring the first SMW.

```
smw2# scp -p user@host:~/cray_bootraid_config.yaml \
/var/adm/cray/cray_bootraid_config.yaml
```

3. Install in bootstrap mode. Specify the storage configuration template (from `smw1`) by using the `--storage-config` parameter with the path to the file on `smw2`.

```
smw2# /media/SMW/SMWinstall --mode bootstrap \
--storage-config /var/adm/cray/cray_bootraid_config.yaml
```

4. Display `cray_bootraid` information.

```
smw# . /opt/modules/default/etc/modules.sh
smw# module use /opt/cray/ari/modulefiles
smw# module load impi
smw# cfgset search -s cray_bootraid -l basic global
smw# cfgset search -s cray_bootraid -l advanced global
```

2.5.2.4 Provision SMW Storage on the Second SMW

About this task

The provision-storage mode of `SMWinstall` uses the boot RAID configuration template (`cray_bootraid_config.yaml`) to provision persistent storage on the boot RAID by creating LVM volume groups and LVM volumes. This is a non-interactive procedure if bootstrap mode has already been completed, which uses the configurator to gather the necessary site-specific configuration information.

Procedure

1. Provision storage for the default SMW storage set. Use this command only if using an SMW storage set called "smwdefault," which is the default.

```
smw2# /media/SMW/SMWinstall --mode=provision-storage
```

2. Wait until `SMWinstall` finishes.

The following warning can be safely ignored if it appears:

```
WARNING: Volume group smw_postgres_vg does not have any devices defined
```

When the provision-storage installer mode completes successfully, the system is ready for the installation of SMW and CLE software.

2.5.2.5 Run the Installer for an Initial Installation

Prerequisites

This procedure assumes that all of the SLES12 ISOs are in `/root/isos`.

About this task

This procedure installs SMW and CLE software together on the second SMW of an SMW HA system to ensure that there is a matched set of software and configuration.

Procedure

1. Set the `SNAPSHOT` environment variable to the name of the snapshot used for the installation of the first SMW (smw1).

Setting a variable here enables better command substitution in later commands dealing with snapshots. The SMW will boot from this snapshot.

IMPORTANT: This snapshot must have exactly the same name as the release snapshot used for the first SMW (smw1), which was saved in [Install the SMWHA Software on the First SMW](#) on page 12 (this is NOT the same as the SNAPSHOT_HA variable, which will be used for the installation of the HA software). Using different snapshots results in HSS database (MySQL) inconsistencies, which causes problems at failover.

```
smw2# export SNAPSHOT=<saved_smw1_snapshot_name>
smw2# echo $SNAPSHOT
```

2. Install SMW and CLE software and security updates together.

It is possible to install both SMW media and CLE media with a single command to create a unified "release" that is tagged as a snapshot on the SMW system. Run the `SMWinstall` program and tell it where the CLE media is. This invocation creates the "target" snapshot, which was named in step 1, and then installs into that target snapshot (note that in the absence of an existing target snapshot, the installer creates one from the current running snapshot by default). The installer assumes that all of the SLES12 ISOs are in `/root/isos`.

IMPORTANT: The SLE media must be specified before the CLE media on the command line so that SUSE security updates are installed before the CLE software is installed.

```
smw2# /media/SMW/SMWinstall \
--plus-media=/root/isos/${SLE_SOFTWARE}.iso \
--plus-media=/root/isos/cle-${CLE_RELEASE}-${CLE_SOFTWARE}.iso \
--target=${SNAPSHOT}
```

It will take about 25 minutes to run a combined installation of SMW, CLE, and security updates for the first time on the SMW. The output of `SMWinstall` provides several command hints, including these three:

- snaputil default** The first command hint (snaputil default) is used to ensure that the SMW is booted from the correct (new) snapshot, which is essential to a successful reboot.
- snaputil chroot** The second command hint (snaputil chroot) is used in the software update process and may be used at other times to look around inside the snapshot.
- snaputil delete** The third command hint (snaputil delete) should be used only if this site needs to remove the newly created snapshot for any reason.

Logs will be in `/var/adm/cray/logs/install` for each invocation of `SMWinstall`.

3. Check new snapshot software versions.

When `SMWinstall` completes, check the snapshot details for the expected SMW and CLE release versions.

```
smw2# /media/SMW/snaputil show ${SNAPSHOT}
active_maps      : None
boot menu       : False
booted          : False
btrfs_object_id : 398
cle_version      : 6.0.2042
created         : 2016-11-03 13:34:24
default         : True
initrd          : initrd-3.12.51-52.39-default
kernel          : vmlinuz-3.12.51-52.39-default
kernels (avail) :
    vmlinuz-3.12.28-4-default
    vmlinuz-3.12.51-52.39-default
name            : smw-8.0.2043_cle-6.0.2042.20161019
parent*        : SLES12.Sep26
path           : /media/root-sv/snapshots/smw-8.0.2043_cle-6.0.2042.20161019
```

```

read-only      : False
smw_version    : 8.0.2043
smwha_version  : None
storage_set    : smwdefault
subvolumes     :
    /var/lib/mysql:smw-8.0.2043_cle-6.0.2042.20161019
    /var/opt/cray/repos:smw-8.0.2043_cle-6.0.2042.20161019
total size     : n/a
unshared size  : n/a
updated        : 2016-11-03 14:08:27.613467

```

NOTE: If this is a subsequent fresh install instead of the very first fresh install, the "parent" entry (indicated by an asterisk in the preceding example), will look like this instead:

```
parent          : SLES12
```

4. Change build_images parameter in `install.cle.conf`.

If the `build_images` parameter was changed prior to running the installer, reset it to its original value so that the `install.cle.conf` files are identical on both SMWs.

```
smw2# grep "build_images" /var/adm/cray/install.cle.conf
```

The SMW is now ready to reboot, which starts with setting the default snapshot to boot from. Trying to boot the SMW without first setting the default snapshot will result in an unbootable SMW.

2.5.2.6 Set Default Snapshot and Boot the SMW

Prerequisites

This procedure assumes that the snapshot variable has been set and the SMW and CLE software has been installed.

About this task

When the `SMWinstall` command was invoked in the previous procedure, it provided several suggested `snaptutil` commands that ensure that the snapshot target is set as the default snapshot for the next boot. This procedure uses one of the commands to ensure that the correct snapshot is used to boot the SMW.

Procedure

1. Set the release snapshot as the default.

IMPORTANT: Do not skip this step. If the SMW is rebooted without first setting the default snapshot, the SMW becomes unbootable.

Fresh install For a fresh install:

```
smw2# /media/SMW/snaptutil default ${SNAPSHOT}
```

Software update For a software update:

```
smw2# snaptutil default ${SNAPSHOT}
```


2. Reboot the SMW to switch to the new release.

```
smw2# reboot
```

2.5.3 Configure the Second SMW for CLE System Hardware

Prerequisites

Before beginning these procedures, the SMW must be booted to a release snapshot.

About this task

Use the following procedures in the order listed to configure the second SMW after installing the SMW and CLE software.

1. Start a Typescript File on the Second/Passive SMW
2. [Change the HSS Data Store \(MariaDB\) Root Password on the Second SMW](#) on page 50
3. [Make a Post-install Snapshot using snaputil](#) on page 51

2.5.3.1 Start a Typescript File on the Second/Passive SMW

About this task

Sites can make as few or as many typescripts as they deem useful. Cray recommends starting a typescript file on each SMW in an SMW HA system at these software installation/configuration milestones:

- just before installing a new software release
- just before configuring the newly installed software

All of these steps apply to the second or 'passive' SMW.

Procedure

1. Log in as root to the SMW.
2. (First time only) Create a release directory for the typescript file.

```
smw2# mkdir -p /var/adm/cray/release
```

3. Change to the release directory.

```
smw2# cd /var/adm/cray/release
```

4. Set a variable equal to today's date.

```
smw2# export TODAY=`date +%Y%m%d`  
smw2# echo $TODAY
```

5. Start a typescript file.

```
smw2# script -af ${TODAY}.suffix
```

For *suffix*, substitute a unique string to distinguish among typescript files, such as `install.1` or `update.2`.

6. Change prompt to include a timestamp.

```
smw2# PS1="\u@\h:\w \t# "
```

2.5.3.2 Change the HSS Data Store (MariaDB) Root Password on the Second SMW

About this task

The HSS data store (MariaDB) root password on the second SMW must match the password on the first SMW. If the password was changed on the first SMW, edit the `/root/.my.cnf` file to change the stored password for the second SMW as well, and then change the password on the second SMW (smw2).

Procedure

1. Edit `/root/.my.cnf` to change the password, substituting the first SMW's MariaDB root password for *MariaDB-password*.

```
smw2# vi /root/.my.cnf
[client]
user=root
password=MariaDB-password
```

If this file does not yet exist, create it and add the lines shown in the example, substituting the new password for the placeholder *MariaDB-password*.

2. Save changes and exit the editor.
3. Ensure that only root can see or write to the `/root/.my.cnf` file.

```
smw2# chmod 600 /root/.my.cnf
```

4. Set or change the MariaDB root password.

```
smw2# mysqladmin -uroot password -p
```

Do one of the following at the prompt:

- To **set** the root password for fresh installs or after the database has been reinitialized, press **Enter** to enter an empty string, the default initial password.

```
Enter password: <cr>
```

- To **change** the root password, enter the existing password.

```
Enter password: existing_password
```

At these prompts, enter the new root password, and then enter it again.

```
New password:
Confirm new password:
```

5. Apply any ansible plays that consume global config set data.

```
smw2# /etc/init.d/cray-ansible start
```

After the SMW HA cluster has been configured, the MySQL database is shared between both SMWs in the HA cluster, so there is no need to edit the `/root/.my.cnf` file again. Once the cluster is fully functional, the administrator can use the `mysqladmin` command on one SMW to reset the MySQL root password.

2.5.3.3 Make a Post-install Snapshot using snaputil

About this task

This procedure uses `snaputil` to make an archival snapshot of the system after installing a new software release (fresh install or software update) and before configuring the SMW for CLE system hardware.

Best Practice. Make a snapshot and back up the config set at the same time to keep them in sync. Name the snapshot and config set backup using the same suffix and date/time stamp to help administrators identify which snapshot and config set backup pairs belong together.

Procedure

1. List the available snapshots on the system.

```
smw2# snaputil list
```

2. Set the `SNAPSHOT` environment variable using the currently booted snapshot name.

```
smw2# export SNAPSHOT=$(snaputil list |grep ^cur| awk '{print $2}')
smw2# echo $SNAPSHOT
```

Setting a variable for the snapshot name enables better command substitution in later commands dealing with snapshots.

This is especially important for SMW HA systems because it makes it easier to use the exact same snapshot name for both SMWs. Using different snapshots results in HSS database (MySQL) inconsistencies, which causes problems at failover.

3. Create a new snapshot.

```
smw2# snaputil create ${SNAPSHOT}.postinstall
```

2.5.4 Install the SMWHA Software on the Second SMW

Prerequisites

This procedure requires the following ISOs:

- `SLE-12-HA-DVD-x86_64-GM-CD1.iso`
- `smwha-sleha12sp0-12.0.3079-201701241124.iso`

Procedure

1. Log in as root to the second SMW.
2. Create a Cray release directory, if necessary, and change to that directory.

```
smw2# mkdir -p /var/adm/cray/release
smw2# cd /var/adm/cray/release
```

3. Copy the SLEHA ISO, `SLE-12-HA-DVD-x86_64-GM-CD1.iso`, to the `/root/isos` directory.
4. Copy the SMWHA release ISO, `smwha-sleha12sp0-12.0.3079-201701241124.iso`, to the `/root/isos` directory.
5. Set the final HA snapshot name on the second SMW. This name **must** match the final HA snapshot name on the first SMW (shown as `saved_snapshot_name`).

```
smw2# export SNAPSHOT_HA=$saved_snapshot_name
```

IMPORTANT: From this point on, the first and second SMW **must** use exactly the same snapshot names.

6. Install the SMWHA software on the second SMW.

```
smw2# mkdir -p /media/SMWHA
smw2# mount -o loop,ro /root/isos/smwha-sleha12sp0-12.0.3079-201701241124.iso \
/media/SMWHA
smw2# /media/SMWHA/SMWHAinstall --target $SNAPSHOT_HA
smw2# snaputil default $SNAPSHOT_HA
```

2.6 Reboot the Second SMW and Power On the First SMW

Procedure

1. Reboot the second SMW and wait for it to reboot completely.

```
smw2# reboot
```

2. After the second SMW is completely rebooted, power up the first SMW and wait for it to reboot completely. This can be done from the iDRAC.

3. After the first SMW is completely rebooted, correct the zypper repo type.

Because `SMWinstall` incorrectly sets the zypper repo type on the second SMW of an SMW HA system during installation, use these two commands on each SMW to correct it.

```
smw1# sed -i 's/type=rpm-md/type=plaindir/' /etc/zypp/repos.d/*.repo
smw1# zypper refresh
```

```
smw2# sed -i 's/type=rpm-md/type=plaindir/' /etc/zypp/repos.d/*.repo
smw2# zypper refresh
```

After the first SMW has successfully booted and the zypper repo type has been corrected, continue to the procedures to configure the SMW HA cluster.

2.7 Configure the SMW HA Cluster

Use the following procedures in the order listed to configure the SMW HA cluster after installing the SMWHA software on the second SMW and completing the procedure to reboot both SMWs.

1. [Gather SMW HA Cluster Information](#) on page 53
2. [Configure Required Cluster Settings](#) on page 55
3. [Add the Second SMW to smw_nodes Node Group](#) on page 64

2.7.1 Gather SMW HA Cluster Information

The SMW HA cluster configuration procedure requires fixed and site-specific IP addresses, host names, ports, and passwords.

Fixed IP Addresses for an SMW HA System

An SMW HA cluster uses the following fixed IP addresses. These IP addresses are set by default and are not site dependent.

Table 2. Fixed IP Addresses for an SMW HA System

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 - Virtual eth1 connection
10.1.1.2	SMW, eth1 - Actual eth1 connection for smw1
10.1.1.3	SMW, eth1 - Actual eth1 connection for smw2
10.2.1.1	SMW, eth2 - Virtual primary heartbeat connection for SMW failover
10.2.1.2	SMW, eth2 - Actual primary heartbeat connection for smw1
10.2.1.3	SMW, eth2 - Actual primary heartbeat connection for smw2
10.2.1.0	Network address to bind to (for eth2 primary heartbeat connection on
10.3.1.1	SMW, eth3 - Virtual eth3 connection
10.3.1.2	SMW, eth3 - Actual eth3 connection for smw1

IP Address	Description
10.3.1.3	SMW, eth3 - Actual eth3 connection for smw2
10.4.1.1	SMW, eth4 - Virtual redundant heartbeat connection for SMW failover
10.4.1.2	SMW, eth4 - Actual redundant heartbeat connection for smw1
10.4.1.3	SMW, eth4 - Actual redundant heartbeat connection for smw2
10.4.1.0	Network address to bind to (for eth4 primary heartbeat connection on smw2)
10.5.1.2	SMW, eth5 - Mirrored PMDB disk connection for smw1
10.5.1.3	SMW, eth5 - Mirrored PMDB disk connection for smw2
127.0.0.1	Localhost (loopback)
225.0.0.1	Multicast IP address for eth4
226.0.0.1	Multicast IP address for eth2

The port used for the heartbeat connection has this default.

Port Number	Description
1694	Multicast port for primary heartbeat connection (for eth2 and eth4 on smw2)

Site-dependent Configuration Values for an SMW HA System

An SMW HA system also requires the following site-dependent host names and IP addresses. If helpful, record the actual values for this site here.

IMPORTANT: The IP addresses for the virtual SMW HA cluster (virtual-smw) and the actual SMWs smw1 and smw2) must be on the same subnet.

Table 3. Site-dependent Configuration Values for an SMW HA System

Description	Example	Actual Value
Virtual host name for SMW HA cluster	virtual-smw	
Host name for first SMW	smw1	
Host name for second SMW	smw2	
iDRAC host name on first SMW	smw1-drac	
iDRAC host name on second SMW	smw2-drac	
Customer network IP address for virtual SMW (the SMW HA cluster)	173.31.73.165	
IP address for first SMW	173.31.73.60	
IP address for second SMW	173.31.73.61	
iDRAC IP address on first SMW	172.31.73.77	

Description	Example	Actual Value
iDRAC IP address on second SMW	172.31.73.79	

Passwords for an SMW HA System

The passwords for an SMW HA system must follow these rules:

- The SMW root password must be the same on each SMW.
- The Integrated Dell™ Remote Access Controller (iDRAC) root password must be the same on each iDRAC.
- The iDRAC root password can be different than the SMW root password.
- The hacluster password on each SMW must be the same as the SMW root password.
- The HA stonith resource passwords must be the same as the iDRAC root password.

This table lists the default values for the passwords that must be the same on both SMWs. Note that the SMW and the iDRAC root passwords have the same default value, but when changed from the default, the SMW root password can be different than the iDRAC root password.

Table 4. Default Passwords for an SMW HA System

ID	Default Password
root on smw1	initial0
root on smw2	initial0
root (iDRAC) on smw1	initial0
root (iDRAC) on smw2	initial0
hacluster (for logging in to crm_gui)	same as SMW root (set during HA configuration)
stonith-1 resource	same as iDRAC root (set during HA configuration)
stonith-2 resource	same as iDRAC root (set during HA configuration)

2.7.2 Configure Required Cluster Settings

Prerequisites

Before beginning this procedure, Cray recommends starting a typescript for each SMW on a local workstation:

```
workstation> script -af my_output_file
Script started, file is my_output_file
workstation> ssh root@smw1
```

Alternatively, create a typescript session in the root home directory and restart the session after the system reboots.

About this task

Use the following procedure to configure the required SMW HA cluster settings. During this procedure, the first SMW (*smw1*) becomes the active SMW. The second (*smw2*) becomes the passive SMW.

Procedure

1. Log into the first SMW (*smw1*) as *root*. Log in directly as *root*; do not use *su* from a different account.

```
workstation> ssh -X root@smw1
```

2. In a separate terminal session, log into the other SMW (*smw2*) as *root*. Log in directly as *root*; do not use *su* from a different account.

```
workstation> ssh -X root@smw2
```

3. Run `check_config` to ensure that both SMWs are running the same SMW and CLE software. Provide the host names and iDRAC IP addresses for both SMWs. The root password is required so that the second SMW can be checked.

```
smw1# check_config smw1 smw2 smw1-iDRAC-IP-addr smw2-iDRAC-IP-addr
Please enter SMW root password and press [ENTER]:
Please enter iDRAC root password and press [ENTER]:
Checking configuration. Please wait...
System is configured correctly! Please check /tmp/ha-check-config.51938.log for
details
```

4. Update the cluster IP addresses.
 - a. Update addresses on *smw1* by running the following command with 0 as the first argument.

```
smw1# /opt/cray/ha-smw/default/hainst/update_addresses 0 smw1 smw2
```

- b. Update addresses on *smw2* by running the following command with 1 as the first argument.

```
smw2# /opt/cray/ha-smw/default/hainst/update_addresses 1 smw1 smw2
```

5. Initialize the cluster on *smw1* with the `ha-cluster-init` command.

```
smw1# ha-cluster-init
```

If the following warning message appears, answer it with 'y' because the HA software does not use SBD.

```
WARNING: No watchdog device found. If SBD is used, the cluster will be unable
to start without a watchdog.
Do you want to continue anyway? [y/N] y
```

Messages continue until a question about overwriting `/root/.ssh/id_rsa`. Answer with 'y' to overwrite this file.


```
Restarting firewall (TCP 30865 5560 7630 21064 open)
Enabling sshd.service
/root/.ssh/id_rsa already exists - overwrite? [y/N] y
```

6. As `ha-cluster-init` runs, it prompts for required information. Answer the following questions to configure the cluster.

IMPORTANT: The SLES defaults are **wrong** for an SMW HA system. Change the default values in this step.

- a. Enter **10.2.1.0** for the network bind address.

Note that the address in brackets is site-dependent, so it will be different for this site, but all sites must set the network bind address to 10.2.1.0.

```
Network address to bind to (e.g.: 192.168.1.0) [172.30.12.0]: 10.2.1.0
```

- b. Enter **226.0.0.1** for the multicast address.

```
Multicast address (e.g.:239.x.x.x): 226.0.0.1
```

- c. Enter **1694** for the multicast port.

```
Multicast port [5405]: 1694
```

- d. Enter **n** (no) for SBD usage.

```
Configure SBD:
...
Do you wish to use SBD? [y/N]: n
WARNING: Not configuring SBD - STONITH will be disabled.
...
Done (log saved to /var/log/sleha-bootstrap.log)
```

- e. Wait for `ha-cluster-init` to finish (normally, about 1 or 2 minutes).

7. Join the second SMW to the cluster.

- a. Execute the `ha-cluster-join` command on `smw2`.

```
smw2# ha-cluster-join
```

- b. If the following warning message appears, answer it with 'y' because the HA software does not use SBD.

```
WARNING: No watchdog device found. If SBD is used, the cluster will be
unable to start without a watchdog.
Do you want to continue anyway? [y/N] y
```

- c. Enter the host name of `smw1`.

```
IP address or hostname of existing node (active SMW): smw1
Enabling sshd.service
```

```
Retrieving SSH keys from smw1
```

- d. If asked for the password, use the `root` password for the SMWs.

```
Password: root-password-for-SMWs
```

- e. Answer the following question with 'y' to overwrite `/root/.ssh/id_rsa`.

```
/root/.ssh/id_rsa already exists - overwrite? [y/N] y
```

8. Check the cluster status to verify that both `smw1` and `smw2` are online.

```
smw1# crm_mon -r1 | grep Online  
Online: [ smw1 smw2 ]
```

The `crm_mon` command displays the SMW host names in alphanumeric order; the first SMW shown is not necessarily the active SMW.

9. Configure `eth4` as the redundant heartbeat channel on `smw1`.

- a. Execute `yast2` to open the YaST2 Control Center.

```
smw1# yast2 cluster
```

For the GUI version of YaST, either execute this command on the SMW console or connect via an `ssh` connection with X11 port forwarding (for example, `ssh -X root@smw1`).

The cluster wizard starts and opens the cluster configuration window.

- b. In the left panel, ensure that Communication Channels is selected.
- c. In the right panel, check the Redundant Channel check box, then enter the following information to configure `eth4` as the redundant channel:
- Bind Network Address: **10.4.1.0**
 - Multicast Address: **225.0.0.1**
 - Multicast Port: **1694**

The screenshot shows the YaST2 Network Configuration window. It has two panels for network configuration. The left panel shows a network interface with 'Bind Network Address' set to 10.2.1.0 and 'Multicast Address' set to 226.0.0.1. The right panel, which is highlighted with a red box, shows a network interface with 'Bind Network Address' set to 10.4.1.0, 'Multicast Address' set to 225.0.0.1, and 'Multicast Port' set to 1694. The 'Redundant Channel' checkbox is checked in the right panel.

IMPORTANT: Be very careful to start the multicast address with 225, even though your fingers want to type the familiar netmask 255. An incorrect multicast address will prevent the cluster from starting.

- d. Double-check your settings. Did you enter 225 (not 255) for the multicast address? Compare the YaST2 screen with the figure above to ensure that the settings are correct.
- e. Click the Finish button.
- f. Close the main YaST2 window to exit `yast2`.

10. Configure `eth4` as the redundant heartbeat channel on `smw2`.

- a. Execute `yast2` to open the YaST2 Control Center.

```
smw2# yast2 cluster
```

For the GUI version of YaST, either execute this command on the SMW console or connect via an `ssh` connection with X11 port forwarding (for example, `ssh -X root@smw2`).

The cluster wizard starts and opens the cluster configuration window.

- b. In the left panel, ensure that **Communication Channels** is selected.
- c. In the right panel, check the Redundant Channel check box, then enter the following information to configure `eth4` as the redundant channel:
 - Bind Network Address: `10.4.1.0`
 - Multicast Address: `225.0.0.1`
 - Multicast Port: `1694`

Transport:
Multicast

Channel
☒ Redundant Channel

Bind Network Address:
10.2.1.0

Multicast Address:
226.0.0.1

Multicast Port:
1694

Member Address:

Bind Network Address:
10.4.1.0

Multicast Address:
225.0.0.1

Multicast Port:
1694

IMPORTANT: Be very careful to start the multicast address with **225**, not 255. An incorrect multicast address will prevent the cluster from starting.

- d. Double-check the settings to verify that **225** was entered (not 255) for the multicast address.
- e. Click the **Finish** button.
- f. Close the main YaST2 window to exit `yast2`.

11. On `smw1`, synchronize the second SMW.

```
smw1# csync2 -xv
```

12. Synchronize the `ssh` host keys. This step makes both SMWs appear to have the same `ssh` host identity when someone connects to the virtual SMW host name or IP address.

- a. On `smw1`, copy the `ssh` host keys to `smw2`.

```
smw1# scp -p /etc/ssh/ssh_host_*key* root@smw2:/etc/ssh
```

- b. On `smw2`, restart the `ssh` daemon and remove any stale keys if they exist.

```
smw2# systemctl restart sshd
smw2# ssh-keygen -R smw2
```

- c. On `smw1`, refresh the `ssh` host keys.

```
smw1# ssh-keygen -R smw2
```

- d. On `smw1`, verify that passwordless `ssh` is still functional to itself and the other SMW. If necessary, answer the prompt or perform the specified action to complete the `ssh` connection.

```
smw1# ssh smw1
...
smw1# ssh smw2
...
```

- e. On `smw2`, verify that passwordless `ssh` is still functional to itself and the other SMW. If necessary, answer the prompt or perform the specified action to complete the `ssh` connection.

```
smw2# ssh smw2
...
smw2# ssh smw1
...
```

13. Reset the login environment on both SMWs by logging out, then logging back in as `root`.

You must log in to the actual (not virtual) SMW as `root`. Do not use `su` from a different account.

In the first terminal window:

```
smw1# exit
workstation> ssh root@smw1
```

In the other terminal window:

```
smw2# exit
workstation> ssh root@smw2
```

14. Configure the site-specific settings in the SMW HA configuration file, `/opt/cray/ha-smw/default/hainst/smwha_args`.

- a. Gather the required host names and IP addresses, as described above, for the following items:

- Virtual host name for the HA cluster
- Virtual IP address for the HA cluster
- iDRAC IP address on the first SMW (called **drac_ip_active** in `smwha_args`)
- iDRAC IP address on the second SMW (called **drac_ip_passive** in `smwha_args`)
- IP address for the second SMW (called **passive_smw_hostname** in `smwha_args`)

- b. Determine the persistent device names for the shared directories on the boot RAID:

- `log_disk_name`
- `db_disk_name`
- `home_disk_name`
- `imps_disk_name`

- c. Edit `/opt/cray/ha-smw/default/hainst/smwha_args`.

```
smw1# vi /opt/cray/ha-smw/default/hainst/smwha_args
```

- d. Replace the following default values with the actual values for the site.

```
--virtual_hostname
cray-smw
--virtual_ip
172.31.73.165
--log_disk_name
/dev/mapper/smw_node_vg-log
--db_disk_name
/dev/mapper/smw_node_vg-db
```

```
--home_disk_name
/dev/mapper/smw_node_vg-home
--imps_disk_name
/dev/mapper/smw_node_vg-imps
--repos_disk_name
/dev/mapper/smw_node_vg-repos
--drac_ip_active
172.31.73.142
--drac_ip_passive
172.31.73.77
--passive_smw_hostname
cray-smw2
--verbose
```

- e. Save your changes and exit the editor.

15. Ensure that /home/crayadm/.gvfs is not mounted.

```
smw1# df -a | grep /home/crayadm/.gvfs && umount -f /home/crayadm/.gvfs
```

16. Ensure that nothing is mounted on /mnt. The SMWHAconfig script uses /mnt to set up the shared storage.

```
smw1# df -a | grep mnt
smw1#
```

17. Configure the SMW HA cluster on the active SMW.

- a. Change to the directory containing the SMWHAconfig command.

```
smw1# cd /opt/cray/ha-smw/default/hainst
```

- b. Execute SMWHAconfig on smw1 only, using the modified configuration file as an argument (prefaced by the @ character).

If necessary, answer a prompt or perform the specified action to complete the ssh connection.

```
smw1# ./SMWHAconfig @smwha_args
2014-08-22 11:1:56,156: INFO      cdir was created
2014-08-22 11:31:56,361: INFO
*****Starting of HA software
installation*****

2014-08-22 11:31:56,361: INFO      cluster virtual IP = 172.31.73.165
2014-08-22 11:31:56,361: INFO      log disk (/var/opt/cray/disk/1) = /dev/
disk/by-id/scsi-360080e500023bfff6000006b1515d9bc9
2014-08-22 11:31:56,361: INFO      db disk (/var/lib/mysql)= /dev/disk/by-id/
scsi-360080e500023bfff6000006b3515d9bdf
2014-08-22 11:31:56,362: INFO      home disk (/home)= /dev/disk/by-id/
scsi-360080e500023bfff6000006b5515d9c01
2014-08-22 11:31:56,362: INFO      verbose mode =
...

```

- c. Enter the root passwords when prompted.

The `SMWHAconfig` command prompts for the SMW and iDRAC root passwords so that it can configure the SMW HA cluster and the iDRAC.

```
Enter SMW root password:
Confirm SMW root password:
Enter current iDRAC root password:
Confirm current iDRAC root password:
Enter new iDRAC root password:
Confirm new iDRAC root password:
```

- d. Wait while `SMWHAconfig` automatically loads the HA cluster configuration settings.
- e. If necessary, examine the log file created by `SMWHAconfig` in `/var/log/SMWHAconfig.log.YYMMDD`. This file is a daily log that is appended to each time `SMWHAconfig` is run during a given day. Ignore the warning message that the CIB has no configuration element, if it appears in the `SMWHAconfig` output.

18. Reboot `smw1` and wait for the reboot to finish.

```
smw1# reboot
```

Before continuing, wait until `smw1` has fully rebooted.

19. Reboot `smw2` and wait for the reboot to finish.

```
smw2# reboot
```

Before continuing, wait until `smw2` has fully rebooted.

20. Take the cluster out of maintenance mode.

```
smw1# maintenance_mode_configure disable
Maintenance mode was disabled
smw1# sleep 300
```

21. Verify cluster status by using either `ha_health` or `crm_mon` to check the status of the HA cluster. Verify that both `smw1` and `smw2` are online and that all resources have started.

- ○ Run `ha_health` periodically until it reports that the cluster is healthy.

```
smw1# ha_health

Cluster State
-----
Health State           : Healthy
Active Node            : smw1
Node-1                  : smw1 (online)
Node-2                  : smw2 (online)
Number of Resources     : 28
Number of Resources Running : 28
Number of Resources Stopped : 0
Maintenance Mode        : disabled
Stonith Mode            : enabled
-----
```

- ○ Use `crm_mon` to ensure that all resources have started.

```
smw1# crm_mon -r1
Last updated: Tue May 17 12:31:43 2016
Last change: Thu May 12 13:53:24 2016
Stack: corosync
Current DC: smw2 (167903491) - partition with quorum
Version: 1.1.12-ad083a8
2 Nodes configured
33 Resources configured

Online: [ smw1 smw2 ]

Full list of resources:

ClusterIP      (ocf::heartbeat:IPaddr2):      Started smw1
ClusterIP1     (ocf::heartbeat:IPaddr2):      Started smw1
ClusterIP2     (ocf::heartbeat:IPaddr2):      Started smw1
ClusterIP3     (ocf::heartbeat:IPaddr2):      Started smw1
ClusterIP4     (ocf::heartbeat:IPaddr2):      Started smw1
ClusterMonitor (ocf::smw:ClusterMonitor):      Started smw1
ClusterTimeSync (ocf::smw:ClusterTimeSync):    Started smw1
.
.
.
```

22. Synchronize the NIMS maps. The link to the active NIMS map must be set so both SMWs have the same setting.

IMPORTANT: This step must be done from `smw1`, that is, the first SMW, which was installed using the procedures in *XC™ Series Software Installation and Configuration Guide (S-2559)*. It is the only SMW that has run `imgbuilder --map`, therefore it is the only SMW that has a current link to the NIMS map. If `smw1` is not already the active SMW at this point, it may be necessary to trigger a failover before running `ha_sync_nims_map`. See "Perform a Manual Failover" in *XC™ Series SMW HA Administration Guide (S-2551)*.

```
smw1# ha_sync_nims_map
```

2.7.3 Add the Second SMW to `smw_nodes` Node Group

About this task

This procedure is necessary to ensure that the `/etc/hosts` file is created with the correct entries.

Procedure

1. Find the host ID of the second SMW.

```
smw2# hostid
```

2. On the first SMW, update the CLE config set (`p0` in this example) to add the host ID of the second SMW to the `smw_nodes` node group.

```
smw1# cfmset update -s cray_node_groups -m interactive p0
```


3. On the first SMW, modify `/etc/hosts` to add the second SMW and the virtual SMW IP addresses.

```
smw1# vi /etc/hosts
172.30.12.90    smw1
172.30.12.43    smw2
172.30.12.172   virtual_smw
```

4. On the first SMW, copy the `/etc/hosts` file to the second SMW.

```
smw1# scp -p /etc/hosts root@smw2:/etc/hosts
```

2.8 Change Default HA Passwords After Installation

About this task

During HA configuration, the passwords for the stonith resources are set to the iDRAC root password. If this site changed the default SMW root and iDRAC root passwords after installing the SMW software, there is no need to change the passwords again. Otherwise, use the following procedure to change the SMW root password and the hacluster and stonith passwords.

To change the iDRAC password, use the procedure in [Change the Default iDRAC Password](#).

The passwords for an SMW HA system must follow these rules:

- The SMW root password must be the same on each SMW.
- The Integrated Dell™ Remote Access Controller (iDRAC) root password must be the same on each iDRAC.
- The iDRAC root password can be different than the SMW root password.
- The hacluster password on each SMW must be the same as the SMW root password.
- The HA stonith resource passwords must be the same as the iDRAC root password.

Procedure

1. Log into the active SMW (for example, `smw1`) as `root`, using the virtual SMW host name (such as `virtual-smw`). After login, the prompt displays the host name of the active SMW.
2. Change the SMW root and hacluster passwords on the active SMW (`smw1`).

The hacluster password must be the same as the SMW root password.

```
smw1# passwd root
smw1# passwd hacluster
```

3. Change the stonith-1 and stonith-2 passwords on the active SMW (`smw1`).

The stonith resource passwords must be the same as the iDRAC root password.

```
smw1# crm resource param stonith-1 set passwd new-password
smw1# crm resource param stonith-2 set passwd new-password
```

4. Change the SMW root and hacluster passwords on the passive SMW (`smw2`), using the same root password as on `smw1`.

The hacluster password must be the same as the SMW root password.

```
smw2# passwd root
smw2# passwd hacluster
```

If the iDRAC root password needs to be changed, proceed to [Change the Default iDRAC Password](#) on page 66. Otherwise, proceed to [Configure Failover Notification](#) on page 66.

2.8.1 Change the Default iDRAC Password

About this task

This procedure describes how to log in to the iDRAC web interface and change a user password.

Procedure

1. Bring up a web browser.
2. Go to: `https://cray-drac`, where `cray-drac` is the name of the iDRAC.
A login screen appears.
3. Log in to the web interface as `root`.
4. Select **iDRAC settings** on the left navigation bar.
5. Expand **iDRAC settings** on the left navigation bar.
6. Select **User Authentication**.
7. Select the user whose password is changing. To change the root password, select `userid 2`.
8. Select **Next**.
9. Select the **Change Password** box and enter the new password in the boxes below it.
10. Select **Apply** to complete the password change.

The password change is complete.

Alternative. Another approach to changing the iDRAC root password is to use `ipmitool` on the SMW command line interface.

```
smw# ipmitool -U root -I lanplus -H <drac-ip-addr> -P <old-drac-password> \
user set password 2 <new-drac-password>
```

2.9 Configure Failover Notification

Prerequisites

Failover notification requires email to be configured on both SMWs. For information about configuring email, see http://www.postfix.org/BASIC_CONFIGURATION_README.html.

About this task

The SMW HA software includes a `Notification` resource that automatically sends email when a failover occurs.

You can configure failover notification either during initial installation or after the HA system is installed and running.

Procedure

1. Execute the `crm resource` command.

```
smw1# crm resource param Notification set email address@thedomain.com
```

NOTE: Only one email address is allowed. To send notifications to multiple addresses, you can create a group email alias that includes these email addresses.

2. Verify the setting.

```
smw1# crm resource param Notification show email  
address@thedomain.com
```

If a failover occurs, the `Notification` resource sends several messages that are similar to the following examples.

```
From: root [mailto:root@smw.none]
Sent: Thursday, June 06, 2013 9:20 PM
To: Cray Cluster Administrator
Subject: ***Alert*** A Failover may have occurred. Please investigate!
Migrating resource
away at Thu Jun 6 21:20:25 CDT 2013 from smw1
```

```
***Alert*** A Failover may have occurred. Please investigate!
Migrating resource away
at Thu Jun 6 21:20:25 CDT 2013 from smw1
```

```
Command line was:
/usr/lib/ocf/resource.d//heartbeat/MailTo stop
From: root [mailto:root@smw.none]
Sent: Thursday, June 06, 2013 9:20 PM
To: Cray Cluster Administrator
Subject: ***Alert*** A Failover may have occurred. Please investigate!
Takeover in progress
at Thu Jun 6 21:20:25 CDT 2013 on smw2
```

```
***Alert*** A Failover may have occurred. Please investigate!
Takeover in progress
at Thu Jun 6 21:20:25 CDT 2013 on smw2
```

```
Command line was:
/usr/lib/ocf/resource.d//heartbeat/MailTo start
```

2.10 Configure the Power Management Database with DRBD for SMW HA

Prerequisites

Before beginning this procedure:

- The SMW HA software must be installed and configured on both SMWs.
- Plan sufficient time for this procedure. Transferring data to a 1 TB disk requires about 10 hours. The SMW HA cluster should be in maintenance mode until the synchronization operation completes. The Cray system (compute and service nodes) can be up and can run jobs during this period.

About this task

The Power Management Database (PMDb) is a PostgreSQL database that contains power management data, event router file system (`erfs`) data, and (optionally) System Environment Data Collections (SEDC) data. The directory `/var/lib/pgsql` is the mount point for the PMDb storage.

On an SMW HA system, the `/var/lib/pgsql` directory is mirrored at a block level to the other SMW as a Distributed Replicated Block Device (DRBD) device. In this configuration, the active SMW mounts `/var/lib/pgsql` and communicates replicated writes over a private TCP/IP connection (`eth5`) to the passive SMW. When a failover occurs, the newly active SMW mounts its local mirrored storage of `/var/lib/pgsql`.

IMPORTANT: DRBD mirroring is required even if a remote PMDb has been configured. Event router and HSS data remains on the DRBD-managed device.

Use this procedure to configure DRBD-mirrored storage on the SMW HA system.



CAUTION: Do not use this procedure if the first SMW's PMDISK has existing data that must be preserved; for example, when converting a non-HA system with a stand-alone SMW to an SMW HA system. To preserve existing data, begin with this procedure instead: [Migrate PostgreSQL Data to DRBD for an SMW HA System](#) on page 113.

Procedure

1. Use two separate terminal sessions for this procedure, with one logged into `smw1` and the other logged into `smw2`.
 - a. Log in to the active SMW (for example, `smw1`) in one terminal session.

```
user@host > ssh root@smw1
smw1#
```

- b. Log in to the other SMW (for example, `smw2`) in a separate terminal session.

```
user@host >ssh root@smw2
smw2#
```

In the following examples, pay attention to the host name in the command prompts to ensure that the commands are executed on the correct SMW.

- Put the cluster into maintenance mode.

```
smw1# maintenance_mode_configure enable
```

- Change the eth5 IP address on smw1.

Edit `/etc/sysconfig/network/ifcfg-eth5` on *smw1* and change `IPADDR` from 10.5.1.1 to 10.5.1.2.

```
smw1# vi /etc/sysconfig/network/ifcfg-eth5
```

The edited file should look similar to this example.

```
BOOTPROTO='static'
IPADDR='10.5.1.2/16'
NAME='eth5 SMW HA DRBD'
PREFIXLEN='16'
STARTMODE='auto'
USERCONTROL='no'
LINK_REQUIRED='no'
```

- Change the eth5 IP address on smw2.

Edit `/etc/sysconfig/network/ifcfg-eth5` on *smw2* and change `IPADDR` from 10.5.1.1 to 10.5.1.3.

```
smw2# vi /etc/sysconfig/network/ifcfg-eth5
```

The edited file should look similar to this example.

```
BOOTPROTO='static'
IPADDR='10.5.1.3/16'
NAME='eth5 SMW HA DRBD'
PREFIXLEN='16'
STARTMODE='auto'
USERCONTROL='no'
LINK_REQUIRED='no'
```

- Reset the eth5 interface on both SMWs.

On smw1:

```
smw1# ifdown eth5; sleep 1; ifup eth5
```

On smw2:

```
smw2# ifdown eth5; sleep 1; ifup eth5
```

- Verify the IP addresses from smw1 by pinging the IP address of eth5 on smw2.

```
smw1# ping -c3 10.5.1.3
```

******* PERFORM STEPS 7 on page 70, 8 on page 70, and 9 on page 70 FOR EACH SMW *******

7. Check that the PMDISK is inserted into the SMW in slot 4 and that the disk has the expected size. A 1TB disk is about 931.5GiB (other disks are much smaller).

In the following command, replace `/dev/disk/by-path/device` with the correct information for the SMW model:

- R815 SMW: `/dev/disk/by-path/pci-0000:05:00.0-sas-0x4433221103000000-lun-0`
- R630 SMW: `/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:4:0`

```
smw# fdisk -l /dev/disk/by-path/device
Disk /dev/disk/by-path/pci-0000:05:00.0-sas-0x4433221103000000-lun-0: 931.5
GiB, 1000204886016 bytes, 1953525168 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disklabel type: dos
Disk identifier: 0x10692081
```

8. Create a new primary partition for PMDISK and write it to the partition table.

In the following command, replace `/dev/disk/by-path/device` with the correct information for the SMW model:

- R815 SMW: `/dev/disk/by-path/pci-0000:05:00.0-sas-0x4433221103000000-lun-0`
- R630 SMW: `/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:4:0`

IMPORTANT: If there are any existing partitions on this disk, manually delete them first using the "d" command in fdisk.

This example shows entering "n" to add a new partition, as a primary partition type, as partition number 1, and accepting the first and last sector so this partition uses all of the space on the disk. Then use "w" to write the new partition table to disk and exit.

```
smw# fdisk /dev/disk/by-path/device
Welcome to fdisk (util-linux 2.25).
Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.

Command (m for help): n
Partition type
  p   primary (0 primary, 0 extended, 4 free)
  e   extended (container for logical partitions)
Select (default p): p
Partition number (1-4, default 1): 1
First sector (2048-1953525167, default 2048): [press return]
Last sector, +sectors or +size{K,M,G,T,P} (2048-1953525167, default
1953525167): [press return]

Created a new partition 1 of type 'Linux' and of size 931.5 GiB.

Command (m for help): w
The partition table has been altered.
Calling ioctl() to re-read partition table.
Syncing disks.
```

9. Verify that the partition on PMDISK has been created.

In the following command, replace `/dev/disk/by-path/partition` with the correct information for the SMW model (the partition name always ends in `-part1`):

- R815
SMW: /dev/disk/by-path/pci-0000:05:00.0-sas-0x4433221103000000-lun-0-part1
- R630 SMW: /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:4:0-part1

```
smw# fdisk -l /dev/disk/by-path/partition
Disk /dev/disk/by-path/pci-0000:05:00.0-sas-0x4433221103000000-lun-0: 931.5
GiB, 1000204886016 bytes, 1953525168 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disklabel type: dos
Disk identifier: 0x10692081

Device                               Boot
Start      End      Sectors  Size Id Type
/dev/disk/by-path/pci-0000:05:00.0-sas-0x4433221103000000-lun-0-part1
2048 1953525167 1953523120 931.5G 83 Linux
```

10. Run the `SMWHAconfig` command on `smw1` to create the DRBD device. Use the `pm_disk_name` option to specify the correct partition name.

In the following command, replace `/dev/disk/by-path/partition` with the correct information for the SMW model (the partition name always ends in `-part1`):

- R815
SMW: /dev/disk/by-path/pci-0000:05:00.0-sas-0x4433221103000000-lun-0-part1
- R630 SMW: /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:4:0-part1

```
smw1# cd /opt/cray/ha-smw/default/hainst
smw1# ./SMWHAconfig --add_disk=pm-fs --device=/dev/drbd_r0 \
--directory=/var/lib/pgsql \
--pm_disk_name=/dev/disk/by-path/partition
```

11. Reboot `smw1` and wait for it to boot completely.

```
smw1# reboot
```

12. Reboot `smw2` and wait for it to boot completely.

```
smw2# reboot
```

13. Take the cluster out of maintenance mode so that `drbd_pgsql` cluster resources start and the DRBD sync will resume.

This step is needed to mount the DRBD device so that the next step affects the top-level directory in that file system rather than the mount point.

```
smw1# maintenance_mode_configure disable
```

When `smw1` and `smw2` rejoin the cluster after rebooting, the primary DRBD disk (in `smw1`) begins to synchronize data to the secondary disk (in `smw2`). DRBD operates at the device level to synchronize the entire contents of the PMDB disk. A full initial synchronization takes a long time, regardless of the size of the PMDB. The time to synchronize a 1 TB external DRBD disk is approximately 10 hours. The Cray system (service and compute nodes) can be booted and can run jobs during this period. Note, however, that the high-speed network (HSN) could experience throttling if there is an SMW failure during the synch operation.

14. Correct the permissions of `/var/lib/pgsql` on the active SMW.

```
smw1# chown postgres:postgres /var/lib/pgsql
smw1# chmod 750 /var/lib/pgsql
```

15. Put the SMW HA cluster back into maintenance mode while waiting for the DRBD sync operation to complete. Cray strongly recommends keeping the SMW HA cluster in maintenance mode to prevent any failover during the sync operation. If a failover were to occur during this period, the newly-active SMW could have an incomplete copy of PMDB data.

- a. Put the SMW HA cluster into maintenance mode on *smw1*.

```
smw1# maintenance_mode_configure enable
```

- b. Check the status of the DRBD sync operation.

```
smw1# cat /proc/drbd
version: 8.4.4 (api:1/proto:86-101)
GIT-hash: 599f286440bd633d15d5ff985204aff4bccffadd build by phil@fat-tyre,
2013-10-11 16:42:48
 0: cs:SyncSource ro:Primary/Secondary ds:UpToDate/Inconsistent C r-----
   ns:695805444 nr:12508 dw:1808112 dr:694131606 al:171 bm:43068 lo:0 pe:2
  ua:0 ap:0 ep:1 wo:f oos:260636656
    [=====>.....] sync'ed: 72.4% (254524/922140)M
   finish: 2:21:07 speed: 30,768 (29,720) K/sec
```

For an explanation of the status information in `/proc/drbd`, see the DRDB User's Guide at [linbit.com: http://drbd.linbit.com/users-guide/ch-admin.html#s-proc-drbd](http://drbd.linbit.com/users-guide/ch-admin.html#s-proc-drbd).

16. When the DRBD sync operation finishes, bring the HA cluster out of maintenance mode on *smw1*.

```
smw1# maintenance_mode_configure disable
```

2.11 Finish Configuring the SMW HA System

About this task

Procedure

1. Bring up the Cray system (service and compute nodes), if not already up.
2. Synchronize ssh user keys between *smw2* and the boot node to enable passwordless access.
 - a. Copy the rsa-key from the first SMW to the second SMW:

```
smw1# scp -pr /root/.ssh/id_rsa* root@smw2:/root/.ssh/
```

- b. Log in to the boot node from the second SMW. Answer reply "yes" when prompted.

```
smw2# ssh boot exit
```

3. If the time zone was changed when installing the base operating system, copy the `localtime` file on the second SMW.

Put the SMW time zone setting where the cabinet and blade controllers can access it. Execute the following command on the second SMW.

```
smw2# cp -p /etc/localtime /opt/tftpboot/localtime
```



CAUTION: If it is necessary to revert to a previous snapshot at some point, use only an HA snapshot — that is, a snapshot created after the SMW HA software was installed and configured. It is dangerous to boot a non-HA snapshot on an HA system because there is a risk of double-mounting the shared file systems, which could cause file system corruption.

3 Update an SMW HA System

Cray provides periodic updates and upgrades to the SMW, CLE, and SMW HA software releases. A software *update* on an SMW HA system involves installing the next minor release for all necessary software components (the SMW, CLE, and SMW HA software). In an update release, the minor version number changes; for example, from SMW SLEHA 11 SP3 to SMW SLEHA 11 SP3 UP01. Update procedures do not change the SMW base operating system. In a software *upgrade*, the major and possibly the minor version numbers change, for example, from SMW 8.0.UP01 to SMW 8.1.UP00.

Follow the procedures in this chapter to update SMW HA system software from CLE 6.0.UP02 / SMW 8.0.UP02 / SMW HA SLEHA12.SP0.UP02 to CLE 6.0.UP03 / SMW 8.0.UP03 / SMW HA SLEHA12.SP0.UP03. The procedures provided here do not change the base operating system version running on the SMW.

update path To use these procedures, this system must be running CLE 6.0.UP02, SMW 8.0.UP02, and SMW HA SLEHA12.SP0.UP02 software, and the SMW must be running the initial release of SUSE Linux Enterprise Server (SLES) version 12.

This section contains the following procedures for updating an SMW HA system:

1. [Prepare for an SMW HA Update](#) on page 74
2. [Update and Configure the XC System on the First SMW](#) on page 78
3. [Update the SMWHA Software on the First SMW](#) on page 78
4. [Update and Configure the XC System on the Second SMW](#) on page 81
5. [Make a Post-update Snapshot using snaputil](#) on page 87
6. [Update the SMWHA Software on the Second SMW](#) on page 88
7. [Reboot the Second SMW after an Update](#) on page 89
8. [Configure the SMW HA Cluster after an Update](#) on page 89

If issues are encountered during the software update, use this procedure to revert to an earlier snapshot, if needed: [Revert to a Previous SMW HA Snapshot](#) on page 93.


3.1 Prepare for an SMW HA Update

In preparation for a software update, do the following:

- Read the *SMW HA Release Errata* and the *SMW HA README* provided with the SMW HA release package for any additional installation-related requirements, corrections to this installation guide, and other relevant information about the release package.
- Read the Field Notices (FN) related to Cray patches to identify any required patches for this release package.
- Read the Field Notices (FN) related to kernel security fixes to identify any changes to this release package. Apply any needed changes before installing the new software.

- Verify that the network connections are in place.
- Know which configuration values are site-specific and which are defaults.
- Be familiar with the default passwords used during the installation process.

When those preparation activities are done, use these procedures prior to beginning the software update.

1.  **CAUTION:** Updating the SMW HA software may overwrite the `/opt/cray/ha-smw/default/hainst/smwha_args` file. If this site wishes to preserve that file for reference, save it now in a location unaffected by the installation.
2. [Start a Typescript File on the First/Active SMW](#) on page 75
3. [Record the Cluster Configuration](#) on page 76
4. Verify that `eth0` and `eth3` are Unmanaged Interfaces
5. [Verify Cluster in a Sane State](#) on page 77
6. [Exit Typescript and Shut Down Second SMW](#) on page 78

3.1.1 Start a Typescript File on the First/Active SMW

About this task

Sites can make as few or as many typescripts as they deem useful. Cray recommends starting a typescript file on each SMW in an SMW HA system at these software installation/configuration milestones:

- just before installing a new software release
- just before configuring the newly installed software

All of these steps apply to the first or 'active' SMW.

Procedure

1. Log in as root to the SMW.
2. (First time only) Create a release directory for the typescript file.

```
smw1# mkdir -p /var/adm/cray/release
```

3. Change to the release directory.

```
smw1# cd /var/adm/cray/release
```

4. Set a variable equal to today's date.

```
smw1# export TODAY=`date +%Y%m%d`
smw1# echo $TODAY
```

5. Start a typescript file.

```
smw1# script -af ${TODAY}.suffix
```

For *suffix*, substitute a unique string to distinguish among typescript files, such as `install.1` or `update.2`.

6. Change prompt to include a timestamp.

```
smw1# PS1="\u@\h:\w \t# "
```

3.1.2 Record the Cluster Configuration

About this task

Display the following information about the current configuration, which will be captured by the typescript started just prior to this procedure.

Procedure

1. Display the virtual hostname.

```
smw# crm resource param fsync show virtual_hostname  
stplabha-smw
```

2. Display the virtual IP address.

```
smw# crm resource param ClusterIP show ip  
172.31.73.165
```

3. Display the iDRAC6 IP addresses.

```
smw# crm resource param stonith-1 show ipaddr  
172.31.73.142  
smw# crm resource param stonith-2 show ipaddr  
172.31.73.77
```

3.1.3 Verify that eth0 and eth3 are Unmanaged Interfaces

About this task

For an SMW HA system, eth0 and eth3 must be configured as unmanaged interfaces in the `cray_global_net` config set worksheet so that Ansible does not manage them. Those configuration settings were set during configuration of the first SMW.

Procedure

1. Verify the configuration settings for eth0 and eth3 before installing or updating the SMW HA software.
2. Change to the global worksheet directory.

```
smw# cd /var/opt/cray/imps/config/sets/global/worksheets
```

3. Search for `customer_ethernet.unmanaged_interface` and ensure that it is set to 'true.'

```
smw# grep "customer_ethernet.unmanaged_interface" cray_global_net_worksheet.yaml
cray_global_net.settings.hosts.data.primary_smw.interfaces.customer_ethernet.unmanaged_interface: true
```

4. Search for **admin_ethernet.unmanaged_interface** and ensure that it is set to 'true.'

```
smw# grep "admin_interface.unmanaged_interface" cray_global_net_worksheet.yaml
cray_global_net.settings.hosts.data.primary_smw.interfaces.admin_interface.unmanaged_interface: true
```

5. If either of these settings is set to 'false,' change it to 'true' using the configurator interactively.

- a. Set the value of the `unmanaged_interface` setting to 'true.'

To make the change, use the "Change a Multival Setting Field during a Configurator Session" procedure (under "Common Tasks When Using the Configurator Interactively") in *XC™ Series Configurator User Guide* (S-2560).

- b. Inspect the contents of the following files on **both** SMWs to see if they have been modified by an Ansible play.

If any of these files have been changed by Ansible, restore them to their original contents.

```
/etc/sysconfig/network/ifcfg-eth0
/etc/sysconfig/network/ifcfg-eth1
/etc/sysconfig/network/ifcfg-eth3
```

Trouble? Contact Cray Customer Support if assistance is needed restoring these files to their original contents.

3.1.4 Verify Cluster in a Sane State

About this task

This procedure looks for specific strings in the output of the `crm_mon` command to verify that the SMW HA cluster is in a sane state before proceeding with the software update.

If any of these checks fail, do not proceed with the update until the problem has been addressed. See the *XC™ Series SMW HA Administration Guide* (S-2551) for help with troubleshooting, or contact Cray Customer Support.

Procedure

1. Check whether both SMWs are online.

```
smw# crm_mon -r1 | grep -i Online
```

Both SMWs should be online.

2. Check whether any resources are stopped.

```
smw# crm_mon -r1 | grep -i Stopped
```

No resources should be stopped.

3. Check for mounted volumes in the SMW volume group.

```
smw# df | grep smw_node_vg
```

At least the following should be mounted from the SMW volume group:

```
/dev/mapper/smw_node_vg-log
/dev/mapper/smw_node_vg-home
/dev/mapper/smw_node_vg-db
/dev/mapper/smw_node_vg-imps
/dev/mapper/smw_node_vg-repos
```

3.1.5 Exit Typescript and Shut Down Second SMW

About this task

These are the final steps in preparing for an SMW HA software update.

Procedure

1. Exit from the typescript session.

```
smw# exit
```

2. Shut down the second SMW before proceeding to update the software on the first SMW.

Note the change in prompt. This command is done on smw2.

```
smw2# shutdown -h now
```

3.2 Update and Configure the XC System on the First SMW

To update and configure the XC system on the first SMW, use chapter 4 of *XC™ Series Software Installation and Configuration Guide* (S-2559). The procedure for the first SMW is the same as for a system with a stand-alone SMW, except for these differences:

- **Second SMW off?** Verify that the second SMW (smw2) has been shut down before proceeding with the SMW/CLE update on the first SMW.
- **Unmanaged interfaces.** During configuration, ensure that eth0 (interface to customer network) and eth3 (interface to admin network) are configured as unmanaged interfaces in the `cray_global_net` configuration template.
- **Patch sets.** Note all patch sets that are applied on the first SMW. The second SMW must have exactly the same patch sets.
- **Customizations.** Record all site customization and local changes that are made on the first SMW. The same changes must be made on the second SMW.

When finished with the update and configuration of the first SMW, return to this guide and continue the SMW HA update process with [Update the SMWHA Software on the First SMW](#) on page 78.

3.3 Update the SMWHA Software on the First SMW

Prerequisites

(REQUIRED) Shut down the Cray system (service and compute nodes) before using this procedure to install and configure the SMWHA software.

This procedure requires the following ISOs:

- SLE-12-HA-DVD-x86_64-GM-CD1.iso
- smwha-sleha12sp0-12.0.3079-201701241124.iso

Procedure

START A TYPESCRIPT FILE

1. Log in as root to the SMW.
2. (First time only) Create a release directory for the typescript file.

```
smw1# mkdir -p /var/adm/cray/release
```

3. Change to the release directory.

```
smw1# cd /var/adm/cray/release
```

4. Set a variable equal to today's date.

```
smw1# export TODAY=`date +%Y%m%d`
smw1# echo $TODAY
```

5. Start a typescript file.

```
smw1# script -af ${TODAY}.suffix
```

For *suffix*, substitute a unique string to distinguish among typescript files, such as `install.1` or `update.2`.

6. Change prompt to include a timestamp.

```
smw1# PS1="\u@\h:\w \t# "
```

COPY ISOs

7. Copy the SLEHA ISO, `SLE-12-HA-DVD-x86_64-GM-CD1.iso`, to the `/root/isos` directory.
8. Copy the SMWHA release ISO, `smwha-sleha12sp0-12.0.3079-201701241124.iso`, to the `/root/isos` directory.

SAVE CURRENT SNAPSHOT NAME

9. Save the name of the current snapshot.

```
smw1# export SNAPSHOT=$(snaputil list |grep ^cur| awk '{print $2}')
```

```
smw1# echo $SNAPSHOT
```

Record this snapshot name offline (not on the SMW) so that it will be accessible later during the installation of the second SMW (smw2) while this SMW (smw1) is powered off.

SET HA SNAPSHOT NAME

10. Set variable for HA snapshot name on first SMW. Select a new target snapshot for the final HA installation target using the currently booted snapshot name.

```
smw1# export SNAPSHOT_HA=$(snaputil list |grep ^cur| awk '{print $2}')-ha
```

11. Record the snapshot name. The exact same snapshot name **must** be used on the second SMW.

IMPORTANT: Both SMWs **must** use the same snapshot name. For SMW HA systems, it is **very important** to use the exact same snapshot name for both SMWs. Failing to do this will result in HSS database (MySQL) inconsistencies between the snapshots.

```
smw1# echo $SNAPSHOT_HA
```

INSTALL HA SOFTWARE

12. Install the SMWHA software on first SMW.

```
smw1# mkdir -p /media/SMWHA
```

```
smw1# mount -o loop,ro /root/isos/smwha-sleha12sp0-12.0.3079-201701241124.iso \
```

```
/media/SMWHA
```

```
smw1# /media/SMWHA/SMWHAinstall --target $SNAPSHOT_HA
```

```
smw1# snaputil default $SNAPSHOT_HA
```

COLLECT INFORMATION FOR THE SECOND SMW

13. Record information and save the files required for the second SMW.

- a. Record the final HA snapshot name used for the first SMW. The second SMW must use the same final snapshot name.
- b. Copy the /var/adm/cray/install.cle.conf file to a remote system.

```
smw1# scp -p /var/adm/cray/install.cle.conf user@host:~/.
```

- c. Record the HSS data store (MariaDB) root password, if changed on the first SMW.
- d. Record any site customization and local changes that were done on the first SMW. These changes must be duplicated exactly on the second SMW.

SHUT DOWN CLE AND POWER OFF THE FIRST SMW

14. Shut down CLE if it is running.

15. Power down the first SMW.

```
smw1# shutdown -h now
```

IMPORTANT: After the first SMW has been powered off, do not turn it back on until directed to do so during the cluster configuration procedure (after all software has been installed on the second SMW).

3.4 Update and Configure the XC System on the Second SMW

To update and configure the second SMW in an SMW HA system, perform these steps/procedures.

1. Verify that the first SMW has completed shutdown and is now powered off.
2. Start a Typescript File on the Second/Passive SMW
3. [Collect Software Media](#) on page 82
4. [Mount Software Media and Prepare install.cle.conf](#) on page 83
5. [Run the Installer for a Software Update](#) on page 84
6. Set Default Snapshot and Boot the SMW

3.4.1 Start a Typescript File on the Second/Passive SMW

About this task

Sites can make as few or as many typescripts as they deem useful. Cray recommends starting a typescript file on each SMW in an SMW HA system at these software installation/configuration milestones:

- just before installing a new software release
- just before configuring the newly installed software

All of these steps apply to the second or 'passive' SMW.

Procedure

1. Log in as root to the SMW.
2. (First time only) Create a release directory for the typescript file.

```
smw2# mkdir -p /var/adm/cray/release
```

3. Change to the release directory.

```
smw2# cd /var/adm/cray/release
```

4. Set a variable equal to today's date.

```
smw2# export TODAY=`date +%Y%m%d`  
smw2# echo $TODAY
```

5. Start a typescript file.

```
smw2# script -af ${TODAY}.suffix
```

For *suffix*, substitute a unique string to distinguish among typescript files, such as `install.1` or `update.2`.

6. Change prompt to include a timestamp.

```
smw2# PS1="\u@\h:\w \t# "
```

3.4.2 Collect Software Media

Prerequisites

This procedure assumes that the base operating system is installed on the SMW and the boot RAID is set up.

About this task

The Cray release distribution media consists of one DVD and several other pieces of media that may be on DVDs or furnished as ISO files. These ISO files are available for download at CrayPort (<https://crayport.cray.com>).

The installer requires several ISO files to be available for setting up and installing packages from SLE repositories. The names of these ISOs are hard-coded in the installer configuration, but the containing directory can be anywhere that makes sense for this site.

IMPORTANT: The default location for these ISO files is `/root/isos`. If that default location is not used for this system, specify the correct location for the ISO files by using the `--iso-dir` argument with the `SMWinstall` command.

Procedure

1. Make a directory on the SMW to hold the ISO files, if one does not already exist.

Instead of placing the ISOs directly in `/root/isos`, use these two commands to place that directory into the `btrfs` subvolume `/var/adm/cray`, which is exempt from snapshots. This prevents the large ISO files from unnecessarily increasing the size of snapshots.

```
smw2# mkdir -p /var/adm/cray/release/isos
smw2# ln -s /var/adm/cray/release/isos /root/isos
```

2. Download the SLES 12 distribution ISOs to the ISO directory on the SMW.

- `SLE-12-Module-Legacy-DVD-x86_64-GM-DVD1.iso`
- `SLE-12-Module-Public-Cloud-DVD-x86_64-GM-DVD1.iso`
- `SLE-12-SDK-DVD-x86_64-GM-DVD1.iso`
- `SLE-12-Server-DVD-x86_64-GM-DVD1.iso`
- `SLE-12-WE-DVD-x86_64-GM-DVD1.iso`

3. Download the CentOS 6.5 distribution ISO (`CentOS-6.5-x86_64-bin-DVD1.iso`) to the ISO directory on the SMW.

4. Download CLE 6.0 and SMW 8.0 ISOs to the ISO directory on the SMW.

- SMW release: `smw-8.0.3075-201701182038.iso`
- CLE release: `cle-6.0.3074-201701182038.iso`

5. Download the SLES 12 security updates ISO (`sleupdate-12sp0+161026-201611021158.iso`) to the ISO directory on the SMW.
6. Make a directory on the SMW to hold any patches that may be available on CrayPort, if one does not already exist.

See the release Errata documents for information about patches.

```
smw2# mkdir -p /var/adm/cray/release/patchsets
```

7. Download patches to the patchset directory on the SMW.

Note that patches are RPMs, not binaries.

3.4.3 Mount Software Media and Prepare `install.cle.conf`

Prerequisites

This procedure assumes that the release software media have been collected and placed in the appropriate directories on the SMW.

About this task

This procedure describes how to mount the SMW media, set environment variables, and copy `install.cle.conf` from the first SMW.

Procedure

MOUNT MEDIA

1. Mount SMW media.

- a. Confirm that this is the right SMW media.

```
smw2# ls -l /root/isos/smw*iso
-rw-r--r-- 1 root root 427184128 Nov  9 10:41 smw-8.0.3075-201701182038.iso
```

- b. Set environment variables for the SMW media.

Use the release string (actually, the build ID) and the date-time stamp for the SMW media as the values for `SMW_RELEASE` and `SMW_SOFTWARE`, as shown in this example.

```
smw2# export SMW_RELEASE=8.0.3075
smw2# echo $SMW_RELEASE

smw2# export SMW_SOFTWARE=201701182038
smw2# echo $SMW_SOFTWARE
```

- c. Mount the SMW release media.

```
smw2# mkdir -p /media/SMW
smw2# mount -o loop,ro /root/isos/smw-${SMW_RELEASE}-${SMW_SOFTWARE}.iso \
/media/SMW
```

2. Check the CLE media.

- a. Confirm that this is the right CLE media.

```
smw2# ls -l /root/isos/cle*iso
-rw-r--r-- 1 root root 1146388480 Nov  9 09:22 cle-6.0.3074-201701182038.iso
```

- b. Set environment variables for the CLE media.

Use the release string and the date-time stamp for the CLE media as the values for CLE_RELEASE and CLE_SOFTWARE, as shown in this example.

```
smw2# export CLE_RELEASE=6.0.3074
smw2# echo $CLE_RELEASE

smw2# export CLE_SOFTWARE=201701182038
smw2# echo $CLE_SOFTWARE
```

3. Set an environment variable for the SLES 12 security updates media.

Use the entire name of the SLES 12 security updates media as the environment variable. This will be used when installing SMW and CLE software and SLES 12 security updates together later in the process.

```
smw2# export SLE_SOFTWARE=sleupdate-12sp0+161026-201611021158
smw2# echo $SLE_SOFTWARE
```

PREPARE THE install.cle.conf FILE

These steps use the configuration file `/var/adm/cray/install.cle.conf` that was saved from the first SMW in [Update the SMWHA Software on the First SMW](#) on page 78.

IMPORTANT: Do not turn on the first SMW at this point. If that file was not saved before turning off the first SMW, skip these two steps. Instead, wait until after the cluster has been configured and both SMWs are up, and then complete this task.

4. Retrieve the `install.cle.conf` from the first SMW.

The `install.cle.conf` file contains configuration that controls the installer's image building behavior.

```
smw2# scp -p user@host:~/install.cle.conf /var/adm/cray/install.cle.conf
```

5. Ensure that image building is disabled.

Images were built as part of the smw1 update process, so if the output of this command shows that `build_images` is set to yes, edit `/var/adm/cray/install.cle.conf` and set `build_images` to no.

```
smw2# grep "build_images" /var/adm/cray/install.cle.conf
build_images: no
```

3.4.4 Run the Installer for a Software Update

Prerequisites

This procedure assumes that all of the SLES 12 ISOs have been downloaded as described in [Collect Software Media](#) on page 82.

About this task

This procedure installs SMW and CLE software together to ensure that there is a matched set of software and configuration.

NOTE: Do NOT run the installer from the `/root/isos` directory. Instead, run it from a directory that is not included in any snapshot, such as `/var/adm/cray/release`.

Procedure

1. Set the `SNAPSHOT` environment variable for the release snapshot to the name of the release snapshot used for the update of the first SMW (`smw1`).

Set a variable for the name of the snapshot that will be used to install and configure the software update. The SMW will boot from this snapshot.

IMPORTANT: This snapshot must have exactly the same name as the release snapshot used for the first SMW (`smw1`), which was saved in [Update the SMWHA Software on the First SMW](#) on page 78 (this is NOT the same as the `SNAPSHOT_HA` variable, which will be used for the installation of the HA software). Using different snapshots results in HSS database (MySQL) inconsistencies, which causes problems at failover.

```
smw2# export SNAPSHOT=<saved_smw1_snapshot_name>
smw2# echo $SNAPSHOT
```

2. Install SMW and CLE software and security updates together into the release snapshot.

The `SMWinstall` program installs SMW media, CLE media, and SLES security updates with a single command to create a unified "release" that is tagged as a snapshot on the SMW system.

IMPORTANT: Specify the SLE media before the CLE media on the command line so that SUSE security updates are installed before the CLE software is installed.

```
smw2# /media/SMW/SMWinstall \
--plus-media=/root/isos/${SLE_SOFTWARE}.iso \
--plus-media=/root/isos/cle-${CLE_RELEASE}-${CLE_SOFTWARE}.iso \
--target=${SNAPSHOT}
```

Note: If running the installer on a system that has more than one SMW storage set, or the SMW storage set does not have the name "smwdefault," then invoke `SMWinstall` with this additional argument: `--storage-set=alternatesmwstorage`.

It will take about 25 minutes to run a combined installation of SMW, CLE, and security updates for the first time on the SMW. The output of `SMWinstall` provides several command hints, including these three:

snaputil default	The first command hint (<code>snaputil default</code>) is used to ensure that the SMW is booted from the correct (new) snapshot, which is essential to a successful reboot.
snaputil chroot	The second command hint (<code>snaputil chroot</code>) is used in the next section to switch to a <code>chroot'd</code> environment. Within the new snapshot, the configurator will be run to make any needed changes in configuration due to new or changed features and also to create new boot images containing the new software.
snaputil delete	The third command hint (<code>snaputil delete</code>) should be used only if this site needs to remove the newly created snapshot for any reason.

Logs will be in `/var/adm/cray/logs/install` for each invocation of `SMWinstall`.

Trouble? If the update needs to be uninstalled, or the system must be downgraded to an earlier release, use the `--forceupdate` flag. This flag bypasses the version validation check and passes the `--force` argument to zypper, which performs a full uninstall+reinstall of every specified package.

```
smw2# /media/SMW/SMWinstall \
--plus-media=/root/isos/${SLE_SOFTWARE}.iso \
--plus-media=/root/isos/cle-${CLE_RELEASE}-${CLE_SOFTWARE}.iso \
--target=${SNAPSHOT} --forceupdate
```

3. Check new snapshot software versions.

When `SMWinstall` completes, check the snapshot details for the expected SMW and CLE release versions.

This example shows output for the SMW 8.0.UP02 / CLE 6.0.UP02 release. The output for the current release will show different release values.

```
smw2# /media/SMW/snaputil show ${SNAPSHOT}
active_maps      :
  p0:/var/opt/cray/imps/config/sets/global/nims/maps/p0
boot menu       : False
booted          : False
btrfs_object_id : 359
cle_version     : 6.0.2042
created         : 2016-10-14 11:41:06
default         : False
initrd          : initrd-3.12.51-52.39-default
kernel          : vmlinuz-3.12.51-52.39-default
kernels (avail) :
  vmlinuz-3.12.28-4-default
  vmlinuz-3.12.51-52.39-default
name            : SMW-8.0UP02_CLE-6.0UP02.20161014
parent          : smw-8.0.96_cle-6.0.96.20161014
path            : /media/root-sv/snapshots/SMW-8.0UP02_CLE-6.0UP02.20161014
read-only       : False
smw_version     : 8.0.2040
smwha_version   : 12.0.48
storage_set     : smwdefault
subvolumes      :
  /var/lib/mysql:SMW-8.0UP02_CLE-6.0UP02.20161014
  /var/opt/cray/repos:SMW-8.0UP02_CLE-6.0UP02.20161014
total size      : n/a
unshared size   : n/a
updated         : 2016-10-14 12:30:36.303019
```

4. Change `build_images` parameter in `install.cle.conf`.

If the `build_images` parameter was changed prior to running the installer, reset it to its original value so that the `install.cle.conf` files are identical on both SMWs.

```
smw2# grep "build_images" /var/adm/cray/install.cle.conf
```

The SMW is now ready to reboot, which starts with setting the default snapshot to boot from. Trying to boot the SMW without first setting the default snapshot could result in an unbootable SMW.

3.4.5 Set Default Snapshot and Boot the SMW during a Software Update

Prerequisites

This procedure assumes that the snapshot variable has been set and the SMW and CLE software has been installed.

About this task

This procedure uses `snaputil` to ensure that the correct snapshot is used to boot the SMW.

Procedure

1. Set the release snapshot as the default.

IMPORTANT: Do not skip this step. If the SMW is rebooted without first setting the default snapshot, the SMW becomes unbootable.

```
smw# snaputil default ${SNAPSHOT}
```

2. Verify that the correct snapshot is the default.

```
smw# snaputil list
```

3. Reboot the SMW to switch to the new release.

```
smw# reboot
```

Trouble? If the reboot fails, try again, and this time check the grub boot screen on the SMW or iDRAC console to make sure the correct snapshot is selected. If necessary, use the arrow key to select the correct snapshot. If the grub boot screen did not automatically select the correct snapshot, update the `/boot/grub2/grub.cfg` file to make this snapshot the default. Wait until the SMW completes the boot, then log in as root and use this command:

```
smw# grub2-mkconfig -o /boot/grub2/grub.cfg
```

3.5 Make a Post-update Snapshot using snaputil

About this task

This procedure uses `snaputil` to make an archival snapshot of the system after updating the XC system on the second SMW in an SMW HA cluster during a software update.

Procedure

1. List the available snapshots on the system.

```
smw2# snaputil list
```

2. Set the `SNAPSHOT` environment variable using the currently booted snapshot name.

```
smw2# export SNAPSHOT=$(snaputil list |grep ^cur| awk '{print $2}')
smw2# echo $SNAPSHOT
```

3. Create a new snapshot.

```
smw2# snaputil create ${SNAPSHOT}.postupdate
```

3.6 Update the SMWHA Software on the Second SMW

Prerequisites

This procedure requires the following ISOs:

- SLE-12-HA-DVD-x86_64-GM-CD1.iso
- smwha-sleha12sp0-12.0.3079-201701241124.iso

Procedure

1. Log in as root to the SMW.
2. Create a Cray release directory, if necessary, and change to that directory.

```
smw2# mkdir -p /var/adm/cray/release
smw2# cd /var/adm/cray/release
```

3. Copy the SLEHA ISO, `SLE-12-HA-DVD-x86_64-GM-CD1.iso`, to the `/root/isos` directory.
4. Copy the SMWHA release ISO, `smwha-sleha12sp0-12.0.3079-201701241124.iso`, to the `/root/isos` directory.
5. Set the final HA snapshot name on the second SMW. This name **must** match the final HA snapshot name on the first SMW (shown as `saved_snapshot_name`).

```
smw2# export SNAPSHOT_HA=$saved_snapshot_name
```

IMPORTANT: From this point on, the first and second SMW **must** use exactly the same snapshot names.

6. Install the SMWHA software on the second SMW.

```
smw2# mkdir -p /media/SMWHA
smw2# mount -o loop,ro /root/isos/smwha-sleha12sp0-12.0.3079-201701241124.iso \
/media/SMWHA
smw2# /media/SMWHA/SMWHAinstall --target $SNAPSHOT_HA
smw2# snaputil default $SNAPSHOT_HA
```


3.7 Reboot the Second SMW after an Update

Procedure

1. Reboot the second SMW.

```
smw2# reboot
```

2. After the second SMW has completed booting, put the SMW HA cluster into maintenance mode.

```
smw2# maintenance_mode_configure enable
```

3. Power up the first SMW and wait for it to reboot completely.

4. After the first SMW is completely rebooted, correct the zypper repo type.

Because `SMWinstall` incorrectly sets the zypper repo type on the second SMW of an SMW HA system during installation, use these two commands on each SMW to correct it.

```
smw1# sed -i 's/type=rpm-md/type=plaindir/' /etc/zypp/repos.d/*.repo
smw1# zypper refresh
```

```
smw2# sed -i 's/type=rpm-md/type=plaindir/' /etc/zypp/repos.d/*.repo
smw2# zypper refresh
```

Trouble? If issues are encountered trying to boot either of the SMWs with the software update, use this procedure to revert to an earlier snapshot, if needed: [Revert to a Previous SMW HA Snapshot](#) on page 93.

After the first SMW has successfully booted and the zypper repo type has been corrected, continue to the procedures to configure the SMW HA cluster.

3.8 Configure the SMW HA Cluster after an Update

Prerequisites

Before beginning this procedure, Cray recommends starting a typescript for each SMW on a local workstation:

```
workstation> script -af my_output_file
Script started, file is my_output_file
workstation> ssh root@smw1
```

Alternatively, create a typescript session in the root home directory and restart the session after the system reboots.

About this task

This procedure configures the required SMW HA cluster settings.

Procedure

1. Log into the first SMW (*smw1*) as *root*.

Log in directly as *root*; do not use *su* from a different account.

```
workstation> ssh -X root@smw1
```

2. In a separate terminal session, log into the other SMW (*smw2*) as *root*.

Log in directly as *root*; do not use *su* from a different account.

```
workstation> ssh -X root@smw2
```

3. Run `check_config` to ensure that both SMWs are running the same SMW and CLE software.

Provide the host names and iDRAC IP addresses for both SMWs. The root password is required so that the second SMW can be checked.

```
smw1# check_config smw1 smw2 smw1-iDRAC-IP-addr smw2-iDRAC-IP-addr
Please enter SMW root password and press [ENTER]:
Please enter iDRAC root password and press [ENTER]:
Checking configuration. Please wait...
System is configured correctly! Please check /tmp/ha-check-config.51938.log for details
```

Ignore any error message that NIMS map links do not match between the two SMWs. NIMS maps will be synchronized later in this procedure.

4. (Optional) Check the SMW HA configuration file, `/opt/cray/ha-smw/default/hainst/smwha_args`.

Edit `/opt/cray/ha-smw/default/hainst/smwha_args` and check to see if the site-specific values are still there. If they are not, and if that file was saved for reference purposes earlier in the process ([Prepare for an SMW HA Update](#) on page 74), then retrieve those values from the saved copy of `smwha_args` and restore them to `/opt/cray/ha-smw/default/hainst/smwha_args`.

5. Ensure that nothing is mounted on `/mnt`.

The `SMWHAconfig` script uses `/mnt` to set up the shared storage.

```
smw1# df -a | grep mnt
smw1#
```

6. Configure the SMW HA cluster on the active SMW.

- a. Change to the directory containing the `SMWHAconfig` command.

```
smw1# cd /opt/cray/ha-smw/default/hainst
```

- b. Execute `SMWHAconfig` on *smw1* only.

If necessary, answer a prompt or perform the specified action to complete the `ssh` connection.

```

smw1# ./SMWHAconfig --update
2016-10-14 19:43:58,104: INFO      cdir was created
2016-10-14 19:43:58,300: INFO      ***** Start of HA software update
*****
2016-10-14 19:43:58,301: INFO      verbose mode = None
cib.delete: cib_update shadow CIB deleted
cib.new: cib_update shadow CIB created
2016-10-14 19:43:59,026: INFO      Enter do_node_disable()
2016-10-14 19:43:59,399: INFO      disabled(standby)
2016-10-14 19:43:59,699: INFO      disabled(standby)
2016-10-14 19:43:59,699: INFO      Enter disp_update
2016-10-14 19:43:59,699: INFO      Backup cluster configuration
...

```

- c. Wait while `SMWHAconfig` automatically loads the HA cluster configuration settings.
- d. If necessary, examine the log file created by `SMWHAconfig` in `/var/log/SMWHAconfig.log.YYMMDD`.
This file is a daily log that is appended to each time `SMWHAconfig` is run during a given day.
Ignore the warning message that the CIB has no configuration element, if it appears in the `SMWHAconfig` output.

7. Reboot `smw1` and wait for the reboot to finish.

```
smw1# reboot
```

Before continuing, wait until `smw1` has fully rebooted.

8. Reboot `smw2` and wait for the reboot to finish.

```
smw2# reboot
```

Before continuing, wait until `smw2` has fully rebooted.

9. Take the cluster out of maintenance mode.

```

smw1# maintenance_mode_configure disable
Maintenance mode was disabled
smw1# sleep 300

```

10. Verify cluster is in a sane state before proceeding.

- a. Check whether both SMWs are online.

```
smw# crm_mon -r1 | grep -i Online
```

Both SMWs should be online.

- b. Check whether any resources are stopped.

```
smw# crm_mon -r1 | grep -i Stopped
```

No resources should be stopped.

- c. Check for mounted volumes in the SMW volume group.

```
smw# df | grep smw_node_vg
```

At least the following should be mounted from the SMW volume group:

```
/dev/mapper/smw_node_vg-log
/dev/mapper/smw_node_vg-home
/dev/mapper/smw_node_vg-db
/dev/mapper/smw_node_vg-imps
/dev/mapper/smw_node_vg-repos
```

11. Verify DRBD is in a sane state before proceeding.

```
smw# cat /proc/drbd
```

The output should show a connect state (cs) of 'Connected' and a disk state (ds) of 'UpToDate/UpToDate.'

12. Verify cluster status using either `ha_health` or `crm_mon`.

Run `ha_health` periodically until it reports that the cluster is healthy.

```
smw1# ha_health
```

```
Cluster State
```

```
-----
Health State           : Healthy
Active Node            : smw1
Node-1                 : smw1 (online)
Node-2                 : smw2 (online)
Number of Resources    : 33
Number of Resources Running : 33
Number of Resources Stopped : 0
Maintenance Mode       : disabled
Stonith Mode           : enabled
-----
```

13. Synchronize the NIMS maps.

The link to the active NIMS map must be set so both SMWs have the same setting.

```
smw1# ha_sync_nims_map
```

14. Verify that the `/etc/hosts` file on each SMW has the correct entries.

a. Verify `/etc/hosts` on the first SMW.

```
smw1# vi /etc/hosts
172.30.12.90    smw1
172.30.12.43    smw2
172.30.12.172   virtual_smw
```

b. Verify `/etc/hosts` on the second SMW.

```
smw2# vi /etc/hosts
172.30.12.90    smw1
172.30.12.43    smw2
172.30.12.172   virtual_smw
```

15. (Optional) Make snapshots post SMW HA cluster configuration.

a. Make a snapshot on the first SMW.

```
smw1# export SNAPSHOT=$(snaputil list |grep ^cur| awk '{print $2}')
```

```
smw1# snaputil create ${SNAPSHOT}.posthaconfig
```

- b. Force a failover to the second SMW.

```
smw1# crm node standby; sleep 300; crm node online
```

- c. Make a snapshot on the second SMW.

```
smw2# export SNAPSHOT=$(snaputil list |grep ^cur| awk '{print $2}')
```

```
smw2# snaputil create ${SNAPSHOT}.posthaconfig
```

- d. Force a failover to return to the first SMW.

```
smw2# crm node standby; sleep 300; crm node online
```

3.9 Revert to a Previous SMW HA Snapshot

Prerequisites

Before beginning this procedure, ensure that the CLE system is shut down.

About this task

This procedure reverts an SMW HA system to a previous snapshot, if necessary.



CAUTION: When reverting to a previous snapshot, use only an HA snapshot — that is, a snapshot created after the SMW HA software was installed and configured. It is dangerous to boot a non-HA snapshot on an HA system because there is a risk of double-mounting the shared file systems, which could cause file system corruption.

Procedure

1. Log in to both SMWs as `root`. This procedure assumes that `smw1` is the active SMW, and `smw2` is the passive SMW. The previous snapshot is shown as `HA-snapshot-name-to-restore`.
2. Set the default snapshot on both SMWs to the name of the HA snapshot to be restored. Do not use a non-HA snapshot.

On `smw1`:

```
smw1# snaputil default HA-snapshot-name-to-restore
```

On `smw2`:

```
smw2# snaputil default HA-snapshot-name-to-restore
```

3. Shut down the passive SMW.

```
smw2# shutdown -h now
```

4. Reboot the active SMW.

```
smw1# reboot
```

Wait several minutes until the active SMW completely reboots and all cluster services have started.

5. Power on the passive SMW.

After the passive SMW has finished rebooting, the SMW HA cluster has been reverted to *HA-snapshot-name-to-restore*.

6. Verify that both SMWs have booted the correct snapshot.

On *smw1*:

```
smw1# snaputil list | grep 'cur,def'cur,def
HA-snapshot-name-to-restore      2016-05-18      07:11:28
```

On *smw2*:

```
smw2# snaputil list | grep 'cur,def'cur,def
HA-snapshot-name-to-restore      2016-05-18      07:11:28
```

4 Customize a Preinstalled SMW HA System

About this task

Cray ships SMW HA systems that are completely installed and configured with Cray-specific host names and IP addresses. To complete the configuration on-site, reconfigure the system with site-specific IP addresses (required), change preassigned default host names (optional), and update the cluster configuration files with the changes.

The customization process updates the IP addresses and host names in the following configuration files:

- `/etc/hosts`
- `/etc/hostname`
- `/etc/csync2/csync2.cfg`
- `/etc/csync2/csync2_cray.cfg`
- `/etc/sysconfig/network/ifcfg-eth0`
- `/etc/sysconfig/network/routes`

Procedure

1. Determine the IP addresses and host names for the SMW HA cluster.

An SMW HA system requires the following site-dependent host names and IP addresses. Use this table to record the actual values for the site.

Table 5. Site-dependent Configuration Values for an SMW HA System

Description	Example	Actual Value
Virtual host name for SMW HA cluster	<code>virtual-smw</code>	
Host name for first SMW	<code>smw1</code>	
Host name for second SMW	<code>smw2</code>	
iDRAC host name on first SMW	<code>smw1-drac</code>	
iDRAC host name on second SMW	<code>smw2-drac</code>	
Customer network IP address for virtual SMW (the SMW HA cluster)	<code>173.31.73.165</code>	
IP address for first SMW	<code>173.31.73.60</code>	
IP address for second SMW	<code>173.31.73.61</code>	

Description	Example	Actual Value
iDRAC IP address on first SMW	172.31.73.77	
iDRAC IP address on second SMW	172.31.73.79	

IMPORTANT: The IP addresses for the virtual SMW HA cluster (`virtual-smw`) and the actual SMWs (`smw1` and `smw2`) must be on the same subnet.

2. Before beginning the site customization, the network administrator or site administrator must assign the IP addresses to the corresponding host names for the SMW HA cluster.
3. If any site-customization and local changes were made on the first SMW before the SMWHA software was installed and configured, duplicate these changes on the second SMW.
4. If any patches were installed on the first SMW, ensure that these patches are also installed on the second SMW.
5. Use the following procedures to make the necessary customizations. Note that these procedures require root privilege.
 1. [Change IP Addresses, Host Names, Name Server, and Route Settings on the First SMW](#) on page 96
 2. [Change the Cluster Configuration on the First SMW](#) on page 98
 3. [Change IP Addresses, Host Names, Name Server, and Route Settings on the Second SMW](#) on page 100
 4. [Finish Customizing a Preinstalled SMW HA System](#) on page 102
 5. [Verify Cluster Status After Customization](#) on page 103
 6. [Change Default SMW, iDRAC, and STONITH Passwords After Customization](#) on page 104
 7. (Optional) Make other changes to the cluster, such as the email address for failover notification, the file synchronization list, or the migration threshold for cluster resources. See [Optional Cluster Configuration Changes](#) on page 106.

4.1 Change IP Addresses, Host Names, Name Server, and Route Settings on the First SMW

Prerequisites

Before beginning this procedure:

- Ensure that the Ethernet cables are connected to the network.
- Verify that the site-specific IP addresses have been assigned to the corresponding host names for the SMW HA cluster.
- Note the IP addresses for the default gateway and name server.
- Ensure that the preinstalled system is backed up.

- Shut down both SMWs, if they are not already down.

About this task

Use the `yast2` utility to customize IP addresses, host names, and other settings on the first SMW.

Procedure

1. Power on the first SMW (`smw1-default`).
2. Log in as `root` on the SMW console. Because this procedure changes host names and IP addresses, you **must** execute this procedure on the SMW console rather than logging in remotely.
3. Execute `yast2` to open the YaST2 Control Center.

```
smw1-default# yast2
```

4. In the right panel, scroll to the Network Devices section and select Network Settings.
5. In the Network Settings window, select the Overview tab.
6. Change the network card setup for the SMW.
 - a. Select `eth0` Customer Network Ethernet, then click the Edit button.
 - b. Enter the IP address of the SMW in the IP Address box.
 - c. Enter the host name of the SMW in the Hostname box.
 - d. Click the Next button to return to the Network Settings window.
7. Define the name servers for the SMW.
 - a. In the Network Settings window, select the Hostname/DNS tab.
 - b. Enter the host name of the SMW in the Hostname box.
 - c. Enter the IP addresses of the name servers into the Name Server boxes. You can define up to three name servers.
 - d. Change the domain name in the Domain Name box to the actual name for the system.
 - e. Change the domain names in the Domain Search box to the actual names for the system.
8. Change the route settings.
 - a. In the Network Settings window, select the Routing tab.
 - b. Enter the IP address for the router in the Default Gateway box.
9. If necessary, change the time zone.
10. To finish the changes, click the OK button. `yast2` writes the configuration changes.
11. Exit `yast2`.

4.2 Change the Cluster Configuration on the First SMW

Prerequisites

Before beginning this procedure:

- Complete the `yast2` changes on the first SMW, as described in [Change IP Addresses, Host Names, Name Server, and Route Settings on the First SMW](#) on page 96.
- Log in as `root` on the first SMW's console. Because this procedure changes host names and IP addresses, this procedure **must** be executed on the SMW console rather than through a remote login.

About this task

This procedure updates the IP addresses and host names in `/etc/csync2/csync2.cfg` and `/etc/csync2/csync2_cray.cfg`. It also updates the CRM cluster configuration file and the DRBD configuration file.

Procedure

1. As `root` on the first SMW, change the synchronization file `/etc/csync2/csync2.cfg`.

- a. Edit `/etc/csync2/csync2.cfg`.
- b. Locate the following lines in the `ha_group` section:

```
host smw1-default
host smw2-default
```

- c. Change these lines to the actual host names for the system, as in this example:

```
host smw1-new
host smw2-new
```

- d. Save the changes and exit the editor.

2. Change the synchronization file `/etc/csync2/csync2_cray.cfg`.

- a. Edit `/etc/csync2/csync2_cray.cfg`.
- b. Locate the following lines in the `cray_group` section:

```
host smw1-default
host smw2-default
```

- c. Change these lines to the actual host names for the system, as in this example:

```
host smw1-new
host smw2-new
```

- d. Locate the following lines in the `user_group` section:

```
host smw1-default
host smw2-default
```

- e. Change these lines to the actual host names for the system, as in this example:

```
host smw1-new
host smw2-new
```

- f. Save the changes and exit the editor.

3. Customize the CRM cluster configuration file.

- a. Edit the cluster configuration file.

```
smw1-default# crm configure edit
```

The configuration file opens in the `vi` editor.

- b. Locate the following lines.

```
node smw1-default \
node smw2-default \
params ip="virtual-smw-default-ip"
params hostname="smw1-default" ipaddr="drac-smw1-ip-default" userid="root"
params hostname="smw2-default" passwd="initial0"
ipaddr="drac-smw2-ip-default"
location stonith-1-loc stonith-1 -inf: smw1-default
location stonith-2-loc stonith-2 -inf: smw2-default
```

- c. Change the host names and IP addresses in these lines to the actual values for the system.

```
node smw1-new \
node smw2-new \
params ip="virtual-smw-new-ip"
params hostname="smw1-new" ipaddr="drac-smw1-ip-new" userid="root"
params hostname="smw2-new" passwd="initial0" ipaddr="drac-smw2-ip-new"
location stonith-1-loc stonith-1 -inf: smw1-new
location stonith-2-loc stonith-2 -inf: smw2-new
```

- d. Save the changes and exit the editor.

4. Customize the DRBD configuration file `/etc/drbd.d/r0.res`.

- a. Edit `/etc/drbd.d/r0.res`.

```
smw1-default# vi /etc/drbd.d/r0.res
```

- b. Enter the actual host names for this system.

Locate the following lines in the file:

```
on smw1-default {
    address 10.5.1.2:7788;
}
on smw2-default {
    address 10.5.1.3:7788;
}
```

Replace `smw1-default` and `smw2-default` with the actual host names for this system, as in this example:

```
on smw1-new {
    address 10.5.1.2:7788;
}
on smw2-new {
    address 10.5.1.3:7788;
}
```

- c. (Optional) Enter the email address that will receive notification of a DRBD split-brain condition.

Locate the following line in the file, and change `root@hostname` to the email address to which any notification of a DRBD split-brain condition should be sent.

```
split-brain "/usr/lib/drbd/notify-split-brain.sh root@hostname";
```

- d. Save changes and exit the editor.

5. Shut down the first SMW. Wait for the system to finish shutting down before continuing to the next procedure.

4.3 Change IP Addresses, Host Names, Name Server, and Route Settings on the Second SMW

Prerequisites

Before beginning this procedure:

- Complete the changes to `csync2` and CRM cluster configuration files on the first SMW, as described in [Change the Cluster Configuration on the First SMW](#) on page 98.
- Shut down the second SMW, if it is not already down.

About this task

Use the `yast2` utility to customize IP addresses, host names, and other settings on the second SMW. It is also necessary to customize the DRBD configuration file here, as was done on the first SMW.

Procedure

1. Power on the second SMW (`smw2-default`).
2. Log in as `root` on the SMW console. Because this procedure changes host names and IP addresses, you **must** execute this procedure on the SMW console rather than logging in remotely.
3. Execute `yast2` to open the YaST2 Control Center.

```
smw2-default# yast2
```

4. In the right panel, scroll to the Network Devices section and select Network Settings.
5. In the Network Settings window, select the Overview tab.
6. Change the network card setup for the SMW.

- a. Select eth0 Customer Network Ethernet, then click the Edit button.
 - b. Enter the IP address of the SMW in the IP Address box.
 - c. Enter the host name of the SMW in the Hostname box.
 - d. Click the Next button to return to the Network Settings window.
7. Define the name servers for the SMW.
- a. In the Network Settings window, select the Hostname/DNS tab.
 - b. Enter the host name of the SMW in the Hostname box.
 - c. Enter the IP addresses of the name servers into the Name Server boxes. You can define up to three name servers.
 - d. Change the domain name in the Domain Name box to the actual name for the system.
 - e. Change the domain names in the Domain Search box to the actual names for the system.
8. Change the route settings.
- a. In the Network Settings window, select the Routing tab.
 - b. Enter the IP address for the router in the Default Gateway box.
9. If necessary, change the time zone.
10. To finish the changes, click the OK button. `yast2` writes the configuration changes.
11. Exit `yast2`.
12. Customize the DRBD configuration file `/etc/drbd.d/r0.res`.
- a. Edit `/etc/drbd.d/r0.res`.

```
smw2-default# vi /etc/drbd.d/r0.res
```

- b. Enter the actual host names for this system.

Locate the following lines in the file:

```
on smw1-default {  
    address 10.5.1.2:7788;  
}  
on smw2-default {  
    address 10.5.1.3:7788;  
}
```

Replace `smw1-default` and `smw2-default` with the actual host names for this system, as in this example:

```
on smw1-new {  
    address 10.5.1.2:7788;  
}  
on smw2-new {  
    address 10.5.1.3:7788;  
}
```

- c. (Optional) Enter the email address that will receive notification of a DRBD split-brain condition.

Locate the following line in the file, and change `root@hostname` to the email address to which any notification of a DRBD split-brain condition should be sent.

```
split-brain "/usr/lib/drbd/notify-split-brain.sh root@hostname";
```

- d. Save changes and exit the editor.

13. Shut down the second SMW. Wait for the system to finish shutting down before continuing to the next procedure.

4.4 Finish Customizing a Preinstalled SMW HA System

Prerequisites

Before beginning this procedure, complete the `yast2` changes on the second SMW, as described in [Change IP Addresses, Host Names, Name Server, and Route Settings on the Second SMW](#) on page 100.

About this task

Update the cluster configuration to finish customizing the SMW HA system.

Procedure

1. Power on the first SMW and wait for it to come up. After the system powers on, the prompt displays the new host name (for example, `smw1-new`).
2. On the first SMW, remove the default host names from the CRM configuration.

In the following commands, replace `smw1-default` with the default (pre-configured) host name of the first SMW. Replace `smw2-default` with the default host name of the second SMW.

```
smw1-new# crm node delete smw1-default
INFO: node smw1-default not found by crm_node
INFO: node smw1-default deleted          =====> deleted

smw1-new# crm node delete smw2-default
INFO: node smw2-default not found by crm_node
INFO: node smw2-default deleted          =====> deleted
```

For each command, ignore the first message that the node is not found. The second message confirms that the node has been deleted.

3. Restart Pacemaker on the first SMW.

```
smw1-new# systemctl restart pacemaker
```

4. Power on the second SMW and wait for it to come up.
After the system powers on, the prompt displays the new host name (for example, `smw2-new`).

5. Copy the synchronization files `/etc/csync2/csync2.cfg` and `/etc/csync2/csync2_cray.cfg` from the first SMW to the second SMW. In the following commands, replace `smw2-new` with the actual host name of the second SMW.

```
smw1-new# scp /etc/csync2/csync2.cfg smw2-new:/etc/csync2/
smw1-new# scp /etc/csync2/csync2_cray.cfg smw2-new:/etc/csync2/
```

6. Synchronize the `csync` files between the first SMW and the second SMW.

```
smw1-new# csync2 -xv
```

If all files are synchronized successfully, `csync2` will finish with no errors.

7. If the time zone was changed on both SMWs, copy the `localtime` file on the second SMW.

Put the SMW time zone setting where the cabinet and blade controllers can access it. Execute the following command on the second SMW.

```
smw2-new# cp -p /etc/localtime /opt/tftpboot/localtime
```

4.5 Verify Cluster Status After Customization

About this task

Ensure that the SMW HA cluster is operating correctly after changing the cluster configuration.

Procedure

1. As `root` on the first SMW, display the cluster status.

```
smw1-new# crm_mon -l
Last updated: Tue May 17 12:31:43 2016
Last change: Thu May 12 13:53:24 2016
Stack: corosync
Current DC: smw2 (167903491) - partition with quorum
Version: 1.1.12-ad083a8
2 Nodes configured
33 Resources configured

Online: [ smw1 smw2 ]

Full list of resources:

ClusterIP      (ocf::heartbeat:IPaddr2):      Started smw1
ClusterIP1     (ocf::heartbeat:IPaddr2):      Started smw1
ClusterIP2     (ocf::heartbeat:IPaddr2):      Started smw1
ClusterIP3     (ocf::heartbeat:IPaddr2):      Started smw1
ClusterIP4     (ocf::heartbeat:IPaddr2):      Started smw1
ClusterMonitor (ocf::smw:ClusterMonitor):      Started smw1
ClusterTimeSync (ocf::smw:ClusterTimeSync):    Started smw1
.
.
.
```

NOTE: `crm_mon` may display different resource names, group names, or resource order on the system.

2. Verify that all resources have started. If necessary, see [Verify the SMW HA Cluster Configuration](#) for additional steps to examine cluster status and fix problems with stopped resources or failed actions.

4.6 Change Default SMW, iDRAC, and STONITH Passwords After Customization

About this task

During HA configuration, the passwords for the stonith resources are set to the iDRAC root password. If this site changed the default SMW root and iDRAC root passwords after installing the SMW software, there is no need to change the passwords again. Otherwise, use the following procedure to change the SMW root password and the hacluster and stonith passwords.

The passwords for an SMW HA system must follow these rules:

- The SMW root password must be the same on each SMW.
- The Integrated Dell™ Remote Access Controller (iDRAC) root password must be the same on each iDRAC.
- The iDRAC root password can be different than the SMW root password.
- The hacluster password on each SMW must be the same as the SMW root password.
- The HA stonith resource passwords must be the same as the iDRAC root password.

Procedure

1. Log into the active SMW (for example, `smw1`) as `root`, using the virtual SMW host name (such as `virtual-smw`). After login, the prompt displays the host name of the active SMW.
2. Change the SMW root and hacluster passwords on the active SMW (`smw1`).

The hacluster password must be the same as the SMW root password.

```
smw1# passwd root
smw1# passwd hacluster
```

3. Change the stonith-1 and stonith-2 passwords on the active SMW (`smw1`).

The stonith resource passwords must be the same as the iDRAC root password.

```
smw1# crm resource param stonith-1 set passwd new-passwd
smw1# crm resource param stonith-2 set passwd new-passwd
```

4. Change the SMW root and hacluster passwords on the passive SMW (`smw2`), using the same root password as on `smw1`.

The hacluster password must be the same as the SMW root password.


```
smw2# passwd root
smw2# passwd hacluster
```

To change the iDRAC password, use the procedure in [Change the Default iDRAC Password](#).

4.6.1 Change the Default iDRAC Password

About this task

This procedure describes how to log in to the iDRAC web interface and change a user password.

Procedure

1. Bring up a web browser.
2. Go to: `https://cray-drac`, where `cray-drac` is the name of the iDRAC.
A login screen appears.
3. Log in to the web interface as `root`.
4. Select **iDRAC settings** on the left navigation bar.
5. Expand **iDRAC settings** on the left navigation bar.
6. Select **User Authentication**.
7. Select the user whose password is changing. To change the root password, select `userid 2`.
8. Select **Next**.
9. Select the **Change Password** box and enter the new password in the boxes below it.
10. Select **Apply** to complete the password change.

The password change is complete.

Alternative. Another approach to changing the iDRAC root password is to use `ipmitool` on the SMW command line interface.

```
smw# ipmitool -U root -I lanplus -H <drac-ip-addr> -P <old-drac-password> \
user set password 2 <new-drac-password>
```

5 Optional Cluster Configuration Changes

IMPORTANT: Before changing the cluster configuration, see [Rules for Changing the SMW HA Cluster Configuration](#) on page 106.

After customizing an SMW HA system, you may choose to make additional configuration changes.

- Change the email address for failover notification.
- Add site-specific files and directories to the synchronization list.
- Change the migration threshold for SMW HA cluster resources.

For more information about making cluster configuration changes, see *XC™ Series SMW HA Administration Guide* (S-2551).

5.1 Rules for Changing the SMW HA Cluster Configuration

The SMW HA system is configured during installation. You can customize the system by changing the failover notification address, resource migration threshold, and list of synchronized files.

When customizing the SMW HA system, follow these rules:

- Do not change the basic cluster configuration, except for the migration threshold (maximum failcount value). You can set the migration threshold for each resource by using the `set_migration_threshold` command.
- Do not attempt to migrate a single resource. All resources must migrate as a group.
- Do not change the system list of synchronized files. You can define which local (site-specific) files are synchronized or excluded from synchronization, but do not add large files or directories to the local list of synchronized files.

5.2 Change Failover Notification

Prerequisites

Failover notification requires email to be configured on both SMWs. For information about configuring email, see http://www.postfix.org/BASIC_CONFIGURATION_README.html.

About this task

The SMW HA software includes a `Notification` resource that automatically sends email when a failover occurs. Failover notification is usually configured during initial installation, but can be changed after the HA system is installed and running.

Procedure

1. Determine the email address for failover notification. Only one email address is allowed. To send notifications to multiple addresses, create a group email alias that includes all necessary email addresses.
2. Execute the following `crm resource command` as `root` on either SMW. Substitute the actual address for `address@thedomain.com`.

```
smw1# crm resource param Notification set email address@thedomain.com
```

3. Verify the setting.

```
smw1# crm resource param Notification show email  
address@thedomain.com
```

5.3 Add Site-specific Files to the Synchronization List

About this task

The file `/etc/csync2/csync2_cray.cfg` specifies the Cray-specific files and directories that must be synchronized, as well as small files that are convenient to keep in sync. For information about the default contents of the synchronization list, see [Synchronized Files](#) on page 108.

IMPORTANT: The `csync2` utility is designed to synchronize small amounts of data. If `csync2` must monitor many directories or synchronize a large amount of data, it can become overloaded and failures may not be readily apparent. For example, do not synchronize the following files or directories:

- `/home`
- `/home/crayadm/.ssh/authorized_keys`
- `/etc/hosts`
- Very large files

Procedure

1. Ensure that the file or directory is small enough for the synchronization list. Cray recommends adding only small files to `/etc/csync2/csync2_cray.cfg`. Use these other methods for large files:
 - Use `scp` to copy a large, static file to the passive SMW, as in this example:


```
smw1# scp -pr /path/file smw2:/path/file
```
 - Use the `rsync` command for directories and files that may change during the copy operation.
2. For each file or directory on the active SMW that you want to synchronize, ensure that the parent directory exists on the passive SMW. In some cases, you must either manually create directories on the passive SMW or copy the directory structure from the active SMW. With either method, be sure that owner, group, and permissions are maintained, because `csync2` can be sensitive to mismatches.
3. Edit the file `/etc/csync2/csync2_cray.cfg` as `root` on the active SMW.

4. To add a file or directory, add the full path (one entry per line) to `/etc/csync2/csync2_cray.cfg`. Comments in this file explain how to make changes.

IMPORTANT: For a symbolic link, only the link itself is synchronized, not the content (destination) of the symbolic link.

5. Save your changes and exit the editor.

The `fsync` resource will synchronize the additional files and directories the next time it runs.

6. If there are local changes to `/etc/hosts` on `smw1`, manually copy `/etc/hosts` to `/etc/hosts` on `smw2`. The customized entries must be above the first section of "XT Cabinet x - y".

```
smw2# cp /etc/hosts /etc/hosts.sav
smw2# scp smw1:/etc/hosts /etc/hosts
```

Then edit the `/etc/hosts` file on `smw2`:

- a. Change IP addresses `10.1.1.x`, `10.2.1.x`, `10.3.1.x`, and `10.4.1.x` to `10.1.1.y`, `10.2.1.y`, `10.3.1.y`, and `10.4.1.y` where if `x` is 2 `y` is 3 and if `x` is 3 `y` is 2.
- b. Change the line `smw1-ip smw1 smw1` to `smw2-ip smw2 smw2`.

5.3.1 Synchronized Files

For files not located on shared storage (boot RAID), the SLEHA Extension software includes the `csync2` utility to synchronize (*sync*) important files between the two SMWs. When a file changes on the active SMW, it is automatically synchronized to the passive SMW. The `csync2` utility synchronizes the required files and directories for the SMW HA cluster, such as `/etc/passwd` and `/opt/cray/hss/*/etc/*`.

File synchronization is automatically configured during initial installation. The file `/etc/csync2/csync2_cray.cfg` lists the Cray-specific files and directories that must be synchronized, as well as small files that are convenient to keep in sync.

File synchronization happens in one direction only: from the active SMW to the passive SMW. If a synchronized file changes on the passive SMW, the change will not be propagated to the active SMW in the course of normal operations and could be overwritten on the passive SMW later if there is a subsequent change to the corresponding file on the active SMW. However, if a failover occurs, the previously passive SMW becomes the active SMW. If the change is still in place, the changed file becomes a candidate for propagation to the other SMW (subject to the rules of file conflict resolution).

The `fsync` resource controls file synchronization operations. Every 100 seconds, `fsync` checks for files that need to be synchronized.

IMPORTANT: If a failover occurs before a file synchronization operation has completed, it could result in the loss of the latest updates.

Very large files are explicitly excluded from synchronization (such as `/opt/cray/hss-images/master`). The `csync2` utility is designed to synchronize small amounts of data. If `csync2` must monitor many directories or synchronize a large amount of data, it can become overloaded and failures may not be readily apparent. Cray recommends that sites do not change the list of synchronized files (or add only small files); instead, copy large files and directories manually to the other SMW.

5.4 Set the Migration Threshold for a Resource

About this task

The `set_migration_threshold` command sets the migration threshold for a resource in an SMW HA cluster. A migration threshold is defined as the maximum number of failures (the failcount) allowed for the resource. If the failcount exceeds this threshold, a failover occurs and management of all cluster resources migrates to the other SMW, making it the active SMW. By default, the migration threshold is 1,000,000.

IMPORTANT: Cray recommends that you either leave migration thresholds at the default values or set them to a very high value until you have experience with SMW HA operation. Migration threshold settings that are too low could cause the resource to be ineligible to run if the failcount exceeds that value on both SMWs. If lower settings are used, Cray recommends that you monitor failcounts regularly for trends and clear the failcount values as appropriate. Otherwise, transient errors over time could push failcount values beyond the migration threshold, which could lead to one of the following scenarios:

- Failovers could be triggered by a transient error condition that might otherwise have been handled by a less disruptive mechanism.
- Failovers might not be possible because both SMWs have exceeded the migration threshold.

NOTE: Execute these commands as `root` on either SMW.

Procedure

1. Determine the resource name.

To display a list of resource names and the status of those resources, use the `crm_resource` command.

```
smw1# crm_resource -l
```

2. Use the `set_migration_threshold` command to change the migration threshold for a resource.

NOTE: For *resource*, specify a particular resource name. For *value*, specify an integer in the range of 0 - 1000000.

```
smw1# set_migration_threshold resource value
```

3. Verify the change.

```
smw1# show_migration_threshold resource
```

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

6 Verify the SMW HA Cluster Configuration

About this task

After rebooting a configured SMW HA system, use this procedure to check that the SMW HA cluster is up and running correctly. After a reboot, wait for 30 to 60 seconds for the cluster to come up fully before beginning this procedure.

Procedure

1. Log in as `root` to the active SMW by using the virtual SMW host name (such as `virtual-smw`). After you have logged in successfully, the prompt displays the host name of the active SMW. The examples in this procedure assume that `smw1` is the active SMW.

```
remote-system% ssh root@virtual-smw
.
.
.
smw1#
```

2. Verify the active SMW by determining where the SMW HA cluster resources are running (such as the `homedir` resource).

```
smw1# crm_mon -r1 | grep homedir
homedir (ocf::heartbeat:Filesystem): Started smw1
```

All resources except `stonith-2` run on the active SMW.

3. Display the cluster status.

```
smw1# crm_mon -r1
Last updated: Tue Jun  7 17:12:32 2016
Last change: Fri Jun  3 16:50:25 2016
Stack: corosync
Current DC: smw2 (167903491) - partition with quorum
Version: 1.1.12-ad083a8
2 Nodes configured
33 Resources configured
```

```
Online: [ smw1 smw2 ]
```

```
Full list of resources:
```

```
ClusterIP      (ocf::heartbeat:IPaddr2):      Started smw1
ClusterIP1     (ocf::heartbeat:IPaddr2):      Started smw1
ClusterIP2     (ocf::heartbeat:IPaddr2):      Started smw1
ClusterIP3     (ocf::heartbeat:IPaddr2):      Started smw1
```

```

ClusterIP4      (ocf::heartbeat:IPaddr2):      Started smw1
ClusterMonitor (ocf::smw:ClusterMonitor):      Started smw1
ClusterTimeSync (ocf::smw:ClusterTimeSync):     Started smw1
HSSDaemonMonitor (ocf::smw:HSSDaemonMonitor):   Started smw1
Notification    (ocf::heartbeat:MailTo):       Started smw1
cray-cfgset-cache (systemd:cray-cfgset-cache):   Started smw1
dhcpd           (systemd:dhcpd.service):        Started smw1
fsync           (ocf::smw:fsync):               Started smw1
hss-daemons    (lsb:rsms):                    Started smw1
repos-fs        (ocf::heartbeat:Filesystem):    Started smw1
stonith-1       (stonith:external/ipmi):        Started smw2
stonith-2       (stonith:external/ipmi):        Started smw1
Resource Group: HSSGroup
  ResourceInit  (ocf::smw:ResourceInit):        Started smw1
  homedir       (ocf::heartbeat:Filesystem):    Started smw1
  md-fs         (ocf::heartbeat:Filesystem):    Started smw1
  pm-fs         (ocf::heartbeat:Filesystem):    Started smw1
  postgresql    (lsb:postgresql):              Started smw1
  mysqld        (ocf::heartbeat:mysql):         Started smw1
Resource Group: IMPSGroup
  imps-fs       (ocf::heartbeat:Filesystem):    Started smw1
  cray-ids-service (systemd:cray-ids-service):  Started smw1
  cray-ansible   (systemd:cray-ansible):        Started smw1
  IMPSystemConfig (ocf::smw:FileSystemConfig):  Started smw1
Resource Group: LOGGroup
  ml-fs         (ocf::heartbeat:Filesystem):    Started smw1
  rsyslog       (systemd:rsyslog.service):      Started smw1
  cray-syslog    (systemd:llmrd.service):       Started smw1
  LOGFileSystemConfig (ocf::smw:FileSystemConfig): Started smw1
ip_drbd_pgsq    (ocf::heartbeat:IPaddr2):      Started smw1
Master/Slave Set: ms_drbd_pgsq [drbd_pgsq]
  Masters: [ smw1 ]
  Slaves: [ smw2 ]

```

Note that `crm_mon` may display different resource names, group names, or resource order on the system.

4. Examine the `crm_mon` output. Verify that each resource has started by looking for `Started smw1` or `Started smw2`. Also look for any failed actions at the end of the output.
5. If not all resources have started or if any failed actions are displayed, execute the `clean_resources` command on either SMW.

IMPORTANT: Before running the `clean_resources` command, log in directly as `root` (instead of using `su` from a `crayadm` login), because `clean_resources` terminates all non-root user sessions.

```

smw1# clean_resources
Cleaning resources on node smw1
Cleaning resource on node=smw1 for resource=ClusterIP
Cleaning up ClusterIP on smw1
Cleaning up ClusterIP on smw2
Waiting for 2 replies from the CRMD.. OK
Cleaning resource on node=smw1 for resource=ClusterIP1
Cleaning up ClusterIP1 on smw1
Cleaning up ClusterIP1 on smw2
Waiting for 2 replies from the CRMD.. OK
.
.
.
Cleaning resource on node=smw2 for resource=ip_drbd_pgsq
Cleaning up ip_drbd_pgsq on smw1

```

```
Cleaning up ip_drbd_pgsql on smw2
Waiting for 2 replies from the CRMd.. OK
Cleaning resource on node=smw2 for resource=drbd_pgsql:0
Cleaning up drbd_pgsql:0 on smw1
Cleaning up drbd_pgsql:0 on smw2
Waiting for 2 replies from the CRMd.. OK
Cleaning resource on node=smw2 for resource=drbd_pgsql:1
Cleaning up drbd_pgsql:1 on smw1
Cleaning up drbd_pgsql:1 on smw2
Waiting for 2 replies from the CRMd.. OK
```

After running `clean_resources`, wait several minutes for cluster activity to settle (check cluster status again with the `crm_mon -r1` command). If the output of this command shows only a subset of the SMW HA services, wait for another minute, then check again. For more information, see the `clean_resources(8)` man page.

7 Additional Procedures for an Installed SMW HA System

The following procedures all assume an installed SMW HA system.

- [Migrate PostgreSQL Data to DRBD for an SMW HA System](#) on page 113

Required only if configuring the PMDISK as a Distributed Replicated Block Device (DRBD) device for the SMW HA system, and the first SMW's PMDISK (`/var/lib/pgsql`) has existing data that must be preserved. If the Power Management Database (PMDb) has already been configured with DRBD, do not use this procedure.

- [Enable Multipath on an Installed SMW HA System](#) on page 119

This procedure assumes that the Cray XC system has already been installed and configured as an SMW HA system without multipath having been enabled. If performing a fresh install, this procedure is not necessary: use the procedures in *XC™ Series Software Installation and Configuration Guide* (S-2559) instead.

- [Re-create Host Certificates to Remedy SSL Certificate Verification Failure](#) on page 123

If SMW HA was recently installed on a system already running DataWarp, that installation may create a new certificate chain, causing SSL certificate verification to fail. Use this procedure to re-create host certificates on the SMW HA system and ensure that DataWarp login nodes are able to verify certificates.

7.1 Migrate PostgreSQL Data to DRBD for an SMW HA System

Prerequisites

IMPORTANT: This procedure is required only if the first SMW's PMDISK (`/var/lib/pgsql`) has existing data that must be preserved before configuring the PMDISK as a Distributed Replicated Block Device (DRBD) device for the SMW HA system. If the Power Management Database (PMDb) has already been configured with DRBD as specified in [Configure the Power Management Database with DRBD for SMW HA](#) on page 68, do not use this procedure.

Before beginning this procedure:

- The SMW HA software must be installed and configured on both SMWs.
- Plan sufficient time for this procedure. Transferring data to a 1 TB disk requires about 10 hours. The SMW HA cluster should be in maintenance mode until the synchronization operation completes. The Cray system (compute and service nodes) can be up and can run jobs during this period.

About this task

The Power Management Database (PMDb) is a PostgreSQL database that contains power management data, event router file system (`erfs`) data, and (optionally) System Environment Data Collections (SEDC) data. The directory `/var/lib/pgsql` is the mount point for the PMDb storage.

On an SMW HA system, the `/var/lib/pgsql` directory is mirrored at a block level to the other SMW as a Distributed Replicated Block Device (DRBD) device. In this configuration, the active SMW mounts `/var/lib/pgsql` and communicates replicated writes over a private TCP/IP connection (`eth5`) to the passive SMW. When a failover occurs, the newly active SMW mounts its local mirrored storage of `/var/lib/pgsql`.

IMPORTANT: DRBD mirroring is required even if a remote PMDb has been configured. Event router and HSS data remains on the DRBD-managed device.

Use this procedure if the first SMW's PMDISK has existing data that must be preserved; for example, when converting a non-HA system with a stand-alone SMW to an SMW HA system. It is not necessary to preserve existing PMDb data. However, consider doing this for the following conditions:

- To migrate the existing PMDb configuration like that set using the `xtpmdbconfig` and `xtpmaction` utilities rather than restarting at default values.
- To migrate existing power, energy, environmental and/or job telemetry.

This procedure backs up existing data, configures DRBD mirrored storage on both SMWs for the PostgreSQL Power Management Database (PMDb), and restores the backed-up data. The data on the `/var/lib/pgsql` file system will be replicated between the two SMWs using DRBD over the `eth5` connection between the two SMWs. This procedure is required even if the system has a remote (off-SMW) PMDb. DRBD mirroring is used for other data on the internal disk, such as event router data.

Procedure

1. Log in to the active SMW as root, using the actual host name of the SMW, not the virtual host name. After logging in, the prompt shows the active SMW's host name (for example, `smw1`).

```
user@host > ssh root@actual_smw_hostname
...
smw1#
```

2. Put the SMW HA cluster into maintenance mode.

Because this procedure shuts down the PostgreSQL server, it is important to put the SMW HA cluster to maintenance mode to prevent unnecessary failovers.

```
smw1# maintenance_mode_configure enable
```

3. Shut down the PostgreSQL database server.

Before copying the contents of the PostgreSQL data directory, the PostgreSQL database server must be cleanly shut down.

```
smw1# systemctl stop postgresql
smw1# systemctl status postgresql
postgresql.service - LSB: Start the PostgreSQL master daemon
Loaded: loaded (/etc/init.d/postgresql)
Active: inactive (dead) since Tue 2016-03-08 17:04:24 CST; 10s ago
```

```
Process: 39912 ExecStop=/etc/init.d/postgresql stop (code=exited, status=0/
SUCCESS)
Process: 22595 ExecStart=/etc/init.d/postgresql start (code=exited, status=0/
SUCCESS)
```

4. Make a file-system-level backup of PMDB data to intermediate storage.

The intermediary storage must be at least as large as the contents of `/var/lib/pgsql`. The following example stores data in the `/root` home directory, though this could just as easily be a remote mounted file system. The key to this archiving step is that permissions are preserved and that the intermediary storage is accessible by the SMW that is active or will be acting as active SMW in the HA cluster.

```
smw1# tar -czpf pmdb_backup.tar.gz -C /var/lib/pgsql .
smw1# ls -l pmdb_backup.tar.gz
-rw-r--r-- 1 root root 6979248 Mar  8 17:19 pmdb_backup.tar.gz
```

5. Unmount `/var/lib/pgsql` (if mounted) and replace with pre-existing backup before using `SMWHAconfig` to set up HA-enabled PMDB storage.

This step restores a previous backup generated by `xtmvpmdb` when the PMDB data was originally moved from the root disk to the dedicated disk. This backup is named `/var/lib/pgsql.MM-DD-YYYYtHH:MM:SS`, where `MM-DD-YYYYtHH:MM:SS` is the timestamp from `xtmvpmdb`. In this example, the backup is named `/var/lib/pgsql.03-07-2016t16:58:30`. This backup is generally small will be used to seed the PMDB when `SMWHAconfig` adds the HA-enabled PMDB storage.

```
smw1# umount /var/lib/pgsql
smw1# rm -rf /var/lib/pgsql
smw1# mv /var/lib/pgsql.03-07-2016t16\:58\:30 /var/lib/pgsql
```

6. Edit `/etc/fstab` to remove or comment out the `/var/lib/pgsql` line.

```
smw1# vi /etc/fstab
...
smw1# grep /var/lib/pgsql /etc/fstab
smw1# echo $?
1
```

7. Log in to the other SMW (for example, `smw2`) in a separate terminal session.

```
user@host >ssh root@smw2
smw2#
```

In the following examples, pay attention to the host name in the command prompts to ensure that the commands are executed on the correct SMW.

8. Change the `eth5` IP address on `smw1`.

Edit `/etc/sysconfig/network/ifcfg-eth5` on `smw1` and change `IPADDR` from `10.5.1.1` to `10.5.1.2`.

```
smw1# vi /etc/sysconfig/network/ifcfg-eth5
```

The edited file should look similar to this example.

```
BOOTPROTO='static'
IPADDR='10.5.1.2/16'
NAME='eth5 SMW HA DRBD'
PREFIXLEN='16'
STARTMODE='auto'
```

```
USERCONTROL='no'
LINK_REQUIRED='no'
```

9. Change the eth5 IP address on smw2.

Edit `/etc/sysconfig/network/ifcfg-eth5` on smw2 and change `IPADDR` from 10.5.1.1 to 10.5.1.3.

```
smw2# vi /etc/sysconfig/network/ifcfg-eth5
```

The edited file should look similar to this example.

```
BOOTPROTO='static'
IPADDR='10.5.1.3/16'
NAME='eth5 SMW HA DRBD'
PREFIXLEN='16'
STARTMODE='auto'
USERCONTROL='no'
LINK_REQUIRED='no'
```

10. Reset the eth5 interface on both SMWs.

On smw1:

```
smw1# ifdown eth5; sleep 1; ifup eth5
```

On smw2:

```
smw2# ifdown eth5; sleep 1; ifup eth5
```

11. Verify the IP addresses from smw1 by pinging the IP address of eth5 on smw2.

```
smw1# ping -c3 10.5.1.3
```

12. Check that the PMDISK is inserted into the SMW in slot 4 and that the disk has the expected size. A 1TB disk is about 931.5GiB (other disks are much smaller).

In the following command, replace `/dev/disk/by-path/device` with the correct information for the SMW model:

- R815 SMW: `/dev/disk/by-path/pci-0000:05:00.0-sas-0x4433221103000000-lun-0`
- R630 SMW: `/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:4:0`

```
smw2# fdisk -l /dev/disk/by-path/device
Disk /dev/disk/by-path/pci-0000:05:00.0-sas-0x4433221103000000-lun-0: 931.5
GiB, 1000204886016 bytes, 1953525168 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disklabel type: dos
Disk identifier: 0x10692081
```

13. Create a new primary partition for PMDISK and write it to the partition table.

In the following command, replace `/dev/disk/by-path/device` with the correct information for the SMW model:

- R815 SMW: `/dev/disk/by-path/pci-0000:05:00.0-sas-0x4433221103000000-lun-0`
- R630 SMW: `/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:4:0`

IMPORTANT: If there are any existing partitions on this disk, manually delete them first using the "d" command in fdisk.

This example shows entering "n" to add a new partition, as a primary partition type, as partition number 1, and accepting the first and last sector so this partition uses all of the space on the disk. Then use "w" to write the new partition table to disk and exit.

```
smw2# fdisk /dev/disk/by-path/device
Welcome to fdisk (util-linux 2.25).
Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.

Command (m for help): n
Partition type
  p   primary (0 primary, 0 extended, 4 free)
  e   extended (container for logical partitions)
Select (default p): p
Partition number (1-4, default 1): 1
First sector (2048-1953525167, default 2048): [press return]
Last sector, +sectors or +size{K,M,G,T,P} (2048-1953525167, default
1953525167): [press return]

Created a new partition 1 of type 'Linux' and of size 931.5 GiB.

Command (m for help): w
The partition table has been altered.
Calling ioctl() to re-read partition table.
Syncing disks.
```

14. Verify that the partition on PMDISK has been created.

In the following command, replace `/dev/disk/by-path/device` with the correct information for the SMW model:

- R815 SMW: `/dev/disk/by-path/pci-0000:05:00.0-sas-0x4433221103000000-lun-0`
- R630 SMW: `/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:4:0`

```
smw2# fdisk -l /dev/disk/by-path/partition
Disk /dev/disk/by-path/pci-0000:05:00.0-sas-0x4433221103000000-lun-0: 931.5
GiB, 1000204886016 bytes, 1953525168 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disklabel type: dos
Disk identifier: 0x10692081

Device                                Start      End      Sectors   Size Id Type          Boot
/dev/disk/by-path/pci-0000:05:00.0-sas-0x4433221103000000-lun-0-part1
2048 1953525167 1953523120 931.5G 83 Linux
```

15. Run the `SMWHAconfig` command on `smw1` to create the DRBD device. Use the `pm_disk_name` option to specify the correct partition name.

In the following command, replace `/dev/disk/by-path/partition` with the correct information for the SMW model (the partition name always ends in `-part1`):

- R815
SMW: /dev/disk/by-path/pci-0000:05:00.0-sas-0x4433221103000000-lun-0-part1
- R630 SMW: /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:4:0-part1

```
smw1# cd /opt/cray/ha-smw/default/hainst
smw1# ./SMWHAconfig --add_disk=pm-fs --device=/dev/drbd_r0 \
--directory=/var/lib/pgsql \
--pm_disk_name=/dev/disk/by-path/partition
```

16. Reboot smw1 and wait for it to boot completely.

```
smw1# reboot
```

17. Reboot smw2 and wait for it to boot completely.

```
smw2# reboot
```

18. Take the cluster out of maintenance mode so that drbd_pgsql cluster resources start and the DRBD sync will resume.

This step is needed to mount the DRBD device so that the next step affects the top-level directory in that file system rather than the mount point.

```
smw1# maintenance_mode_configure disable
```

When *smw1* and *smw2* rejoin the cluster after rebooting, the primary DRBD disk (in *smw1*) begins to synchronize data to the secondary disk (in *smw2*). DRBD operates at the device level to synchronize the entire contents of the PMDB disk. A full initial synchronization takes a long time, regardless of the size of the PMDB. The time to synchronize a 1 TB external DRBD disk is approximately 10 hours. The Cray system (service and compute nodes) can be booted and can run jobs during this period. Note, however, that the high-speed network (HSN) could experience throttling if there is an SMW failure during the synch operation.

19. Correct the permissions of /var/lib/pgsql on the active SMW.

```
smw1# chown postgres:postgres /var/lib/pgsql
smw1# chmod 750 /var/lib/pgsql
```

20. Put the SMW HA cluster back into maintenance mode while waiting for the DRBD sync operation to complete. Cray strongly recommends keeping the SMW HA cluster in maintenance mode to prevent any failover during the sync operation. If a failover were to occur during this period, the newly-active SMW could have an incomplete copy of PMDB data.

- a. Put the SMW HA cluster into maintenance mode on *smw1*.

```
smw1# maintenance_mode_configure enable
```

- b. Check the status of the DRBD sync operation.

```
smw1# cat /proc/drbd
version: 8.4.4 (api:1/proto:86-101)
GIT-hash: 599f286440bd633d15d5ff985204aff4bccffadd build by phil@fat-tyre,
2013-10-11 16:42:48
 0: cs:SyncSource ro:Primary/Secondary ds:UpToDate/Inconsistent C r-----
    ns:695805444 nr:12508 dw:1808112 dr:694131606 al:171 bm:43068 lo:0 pe:2
ua:0 ap:0 ep:1 wo:f oos:260636656
    [=====>.....] sync'ed: 72.4% (254524/922140)M
    finish: 2:21:07 speed: 30,768 (29,720) K/sec
```

For an explanation of the status information in `/proc/drbd`, see the DRDB User's Guide at [linbit.com: `http://drbd.linbit.com/users-guide/ch-admin.html#s-proc-drbd`](http://drbd.linbit.com/users-guide/ch-admin.html#s-proc-drbd).

21. When the DRBD sync operation finishes, bring the HA cluster out of maintenance mode on `smw1`.

```
smw1# maintenance_mode_configure disable
```

22. Restore the pre-existing PMDB data.

TIP: The pre-existing PMDB data can be restored from the backup archive during or after the DRBD sync operation.

- a. Ensure that the PostgreSQL database server is not running.

```
smw1# systemctl stop postgresql
```

- b. Clear out the existing contents of `/var/lib/pgsql`.

```
smw1# rm -rf /var/lib/pgsql/*
```

- c. Restore the backup archive.

```
smw1# tar -xzf pmdb_backup.tar.gz -C /var/lib/pgsql
```

- d. Ensure the correct permissions for `/var/lib/pgsql`.

```
smw1# chown postgres:postgres /var/lib/pgsql
smw1# chmod 750 /var/lib/pgsql
```

23. Start the PostgreSQL server.

```
smw1# systemctl start postgresql
```

7.2 Enable Multipath on an Installed SMW HA System

Prerequisites

This procedure assumes that the Cray XC system has already been installed and configured as an SMW HA system without multipath having been enabled. If performing a fresh install, this procedure is not necessary: use the procedures in XC™ Series Software Installation and Configuration Guide instead.

About this task

This procedure describes how to enable multipath on a Cray XC system that has already been installed and configured as an SMW HA system. Note that multipath does NOT need to be fully cabled to be used. The multipath driver can handle using one path or many.

IMPORTANT: If this system has partitions, repeat any steps that modify 'p0' for each partition. Multipath must be enabled everywhere or nowhere; enabling it on only part of the system causes problems.

Procedure

1. Enable and start the multipath daemon now on each SMW in the HA cluster (active and passive).

```
smw1# systemctl enable multipathd
smw1# systemctl start multipathd

smw2# systemctl enable multipathd
smw2# systemctl start multipathd
```

2. Obtain the host ID of each SMW in the HA cluster and the cnames of any nodes in the system that are connected to the boot RAID with an HBA (host bus adapter).

The system should be bounced or booted for `xtcheckhss` to return a proper list. Run `hostid` for each SMW in the SMW HA system.

```
smw1# hostid
{8 digit hostid}
smw1# xtcheckhss --detail=f --pci

smw2# hostid
{8 digit hostid}
smw2# xtcheckhss --detail=f --pci
```

Look for cnames with HBAs like 'QLogic_ISP2532_8Gb_Fibre_Channel_HBA.'

3. Use the configurator to enable and customize multipath in the global config set on the active SMW.

```
smw1# cfgset update -s cray_multipath -m interactive global
```

- a. Enable multipath.

Enter **E** at the configurator prompt to toggle the enable status of the multipath service, which is disabled by default.

```
Cray Multipath Configuration Service Menu [default: save & exit - Q] $ E
```

- b. Add the host IDs and cnames obtained in step 2 on page 120.

Enter values at the configurator prompt as shown. When prompted to add `node_list` entries, add the host IDs and cnames obtained in step 2 on page 120, one per line. When finished, press **Ctrl-d**. Remember to add the `hostid` of both SMWs.

```
Cray Multipath Configuration Service Menu [default: save & exit - Q] $ 1
...
Cray Multipath Configuration Service Menu [default: configure - C] $ C
...
cray_multipath.settings.multipath.data.node_list
[<cr>=set 0 entries, +=add an entry, ?=help, @=less] $ +
Add node_list (Ctrl-d to exit) $
```

4. Use the configurator to update `cray_bootraid` in the global config set on the active SMW.


```
smw1# cfgset update -s cray_bootraid -m interactive global
```

- a. Select the storage sets setting to configure it.

```
Boot RAID Configuration Service Menu [default: save & exit - Q] $ 1
...
Boot RAID Configuration Service Menu [default: configure - C] $ C
```

- b. For each device in the cledefault and smwdefault storage sets, modify the path name from scsi to dm-uuid-mpath.

This example shows selecting the cledefault (1) volume group (a) boot_node_vg (1) devices (b) field. The * indicates that the selection is to be edited.

```
cray_bootraid.settings.storage_sets
[<cr>=set 2 entries, +=add an entry, ?=help, @=less] $ 1a1b*
```

Remove the "scsi path name, and replace it with the dm-uuid-mpath name.

```
cray_bootraid.settings.storage_sets.data.cledefault.volume_groups.boot_node_vg.devices
[<cr>=set 1 entries, +=add an entry, ?=help, @=less] $ 1-

cray_bootraid.settings.storage_sets.data.cledefault.volume_groups.boot_node_vg.devices
[<cr>=set 0 entries, +=add an entry, ?=help, @=less] $ +
Add devices (Ctrl-d to exit) $ /dev/disk/by-id/dm-uuid-mpath-3600a0980009ec0750000010a5762af70
Add devices (Ctrl-d to exit) $ <Ctrl-d>
```

Set the entries for the boot_node_vg volume group.

```
cray_bootraid.settings.storage_sets.data.cledefault.volume_groups.boot_node_vg.devices
[<cr>=set 1 entries, +=add an entry, ?=help, @=less] $ <cr>
```

Repeat substep b for each device in the cledefault and smwdefault storage sets. Enter * at the prompt to see all storage set entries.

- To select the next cledefault volume group device (sdb_node_vg), enter **1a2b*** at the prompt. If there are more cledefault volume groups, increment the third character to select each one (**1a3b***, **1a4b***, and so forth).
- To select the first smwdefault volume group device (smw_node_vg), enter **2a1b*** at the prompt. If there are more smwdefault volume groups, increment the third character to select each one (**2a2b***, **2a3b***, and so forth).

- c. Set the storage set entries, then save and exit the configurator.

```
cray_bootraid.settings.storage_sets
[<cr>=set 2 entries, +=add an entry, ?=help, @=less] $ <cr>
...
Boot RAID Configuration Service Menu [default: save & exit - Q] $ Q
```

5. Use the configurator to set up inheritance for multipath in the CLE config set of the active SMW.

This example uses 'p0' as the name of the CLE config set. Substitute the actual name used at this site.

```
smw1# cfgset update -s cray_multipath -m interactive p0
```

Enter **I** at the configurator prompt to toggle the inherit status of the multipath service, which is disabled by default. This means that multipath settings in the global config set will be used instead of multipath settings in the CLE config set.

```
Cray Multipath Configuration Service Menu [default: save & exit - Q] $ I
```

Repeat this step for each CLE config set.

6. Validate the config sets and run cray-ansible on the active SMW to apply the config set changes.

- a. Validate the config sets.

```
smw1# cfgset validate global
```

```
smw1# cfgset validate p0
```

- b. Run cray-ansible.

```
smw1# /etc/init.d/cray-ansible start
```

7. For systems using direct-attached Lustre (DAL), update the `dal.fs_defs` file.

Repeat for each partition.

- a. Locate the current `fs_defs` files (typically stored in `/home/crayadm`).

```
smw1# find /home/crayadm -name "*fs_defs*"
```

- b. Find the `fs_defs` files that are currently installed and compare with the one found in `/home/crayadm`.

```
smw1# find /var/opt/cray/imps/config/sets/p0 -name "*fs_defs*"
```

```
smw1# diff /home/crayadm/dal.fs_defs /var/opt/cray/\  
imps/config/sets/p0/lustre/.lctrl/dal.fs_defs.20160205.1454685527
```

- c. Edit the `dal.fs_defs` file to ensure that it has the proper mpath paths in it.

```
smw1# cd /home/crayadm
```

```
smw1# sed -i.nompath 's/\/dev\/disk\/by-id\/scsi\/dev\/disk\/by-id\/dm-uuid-  
mpath/g' \  
dal.fs_defs
```

```
smw1# cp -p dal.fs_defs dal.fs_defs.mpath
```

- d. Install the new `dal.fs_defs` file using `lustre_control`.

```
smw1# lustre_control install -c p0 /home/crayadm/dal.fs_defs
```

8. Shut down all partitions of the Cray system (service and compute nodes).

9. Run `cray-ansible` on the active SMW.

```
smw1# /etc/init.d/cray-ansible start
```

10. Check whether `/etc/lvm/lvm.conf` and `/etc/multipath.conf` were synced to the passive SMW.

They should be synced automatically, but if they are not the same, `scp` both files from the active SMW to the passive SMW.

On the active SMW:

```
smw1# stat /etc/lvm/lvm.conf
smw1# stat /etc/multipath.conf
```

On the passive SMW (smw2 in the example):

```
smw2# stat /etc/lvm/lvm.conf
smw2# stat /etc/multipath.conf
```

11. Enable multipath on the passive SMW.

```
smw2# systemctl enable multipathd
```

12. Put the SMW HA system into maintenance mode.

```
smw1# maintenance_mode_configure enable
```

13. Reboot both SMWs at the same time.

```
smw1# reboot
smw2# reboot
```

14. After both SMWs have completed booting, disable maintenance mode and check cluster status.

```
smw1# maintenance_mode_configure disable
smw1# sleep 300
smw1# crm_mon -r1
```

15. Boot the Cray system.

7.3 Re-create Host Certificates to Remedy SSL Certificate Verification Failure

About this task

Failure to verify an SSL certificate can cause the DataWarp `dwcli` and `dwstat` commands to fail on DataWarp API nodes. The symptom looks like this:

```
login# dwstat all
Connecting to https://c1-0c0s0n2:81 yielded fatal error:
[SSL: CERTIFICATE_VERIFY_FAILED] certificate verify failed (_ssl.c:581)
```

One reason SSL certificate verification may fail is if SMW HA was recently installed on a system already running DataWarp. The installation creates a new certificate chain, thereby invalidating any client certificates that were generated by the prior non-HA installation.

Use this procedure to re-create host certificates on the SMW HA system and ensure that DataWarp login nodes are able to verify certificates.

Procedure

1. Find out which server nodes are configured as DataWarp API nodes.

Cat the following file on any login node:

```
login# cat /etc/opt/cray/dws/dwrest_gw.conf
https://c1-0c0s0n2:81
```

2. On both SMWs (active and passive), remove or move the associated key and crt files for those nodes from the certificate authority.

```
smw# mv /var/opt/cray/certificate_authority/hosts/c1-0c0s0n2.key \
/var/opt/cray/certificate_authority/hosts/c1-0c0s0n2.key.old

smw# mv /var/opt/cray/certificate_authority/hosts/c1-0c0s0n2.crt \
/var/opt/cray/certificate_authority/hosts/c1-0c0s0n2.crt.old
```

3. Run `xtmake_ca` on the active SMW (smw1 in example) to create a new signed host certificate (certificate and key files), using the current certificate authority, for each host found in step 1.

```
smw1# xtmake_ca create c1-0c0s0n2 hosts
```

4. Update and validate the default config set (p0 in this example) on the active SMW.

```
smw1# cfgset update p0

smw# cfgset validate p0
```

5. Run `cray_ansible` on the active SMW to apply the changes there.

```
smw1# /etc/init.d/cray-ansible start
```

6. Refresh the config set on each affected login node.

```
login# /opt/cray/imps-distribution/default/bin/refresh.py
```

7. Run `cray_ansible` on each affected login node to apply the changes there.

```
login# /etc/init.d/cray-ansible start
```

8. Restart `nginx` to pick up the new SSL certificates.

```
login# systemctl restart nginx
```