



Sonexion 1.5 Upgrade Guide

Contents

| | |
|---|----|
| About <i>Sonexion 1.5 Upgrade Guide</i> | 3 |
| Sonexion 1.5.0 Upgrade Introduction..... | 4 |
| Prerequisites..... | 8 |
| Begin Upgrade..... | 10 |
| Upgrade Sonexion 1600/900 Software..... | 17 |
| Post-Installation Steps..... | 28 |
| Troubleshoot from Error Messages..... | 33 |

About Sonexion 1.5 Upgrade Guide

This document provides instructions to upgrade Sonexion systems running software release 1.4.0 to the 1.5.0 software platform.

Audience

This publication is intended for use by Cray field technicians who are trained with Sonexion and familiar with the software upgrade process.

Typographic Conventions

| | |
|---------------------------|---|
| Monospace | A Monospace font indicates program code, reserved words or library functions, screen output, file names, path names, and other software constructs |
| Monospaced Bold | A bold monospace font indicates commands that must be entered on a command line. |
| <i>Oblique or Italics</i> | An <i>oblique</i> or <i>italics</i> font indicates user-supplied values for options in the syntax definitions |
| Proportional Bold | A proportional bold font indicates a user interface control, window name, or graphical user interface button or control. |
| Alt-Ctrl-f | Monospaced hyphenated text typically indicates a keyboard combination |

Record of Revision

| Publication Number | Date | Description |
|--------------------|---------------|-------------------|
| HR5-6099-0 | August 2012 | Original Printing |
| HR5-6099-A | November 2012 | Release 1.1 |
| HR5-6119-A | April 2013 | Release 1.2 |
| S-2532-123 | June 2013 | Release 1.2.3 |
| S-2532-131 | April 2014 | Release 1.3.1 |
| S-2532-14 | December 2014 | Release 1.4 |
| S-2532-15 | June 2015 | Release 1.5 |

Sonexion 1.5.0 Upgrade Introduction

The Sonexion 1.5.0 software upgrade procedure includes steps to do the following:

- Log in to a terminal screen to start the upgrade
- Check for available space
- Create the software upgrade image for a USB flash drive (8GB or larger)
- Run the software upgrade script

Expert level personnel: Technicians or administrators who are authorized from each development or test group will be required to properly install this software upgrade.

Service Interruption Level

This procedure has a service level of Interrupt: the procedure requires taking the Lustre file system offline.

Estimated Service Interval

The time required to upgrade system software depends on several factors, including the specific upgrade path and the technician's familiarity with the upgrade process. The following table shows estimates of the time required to complete individual upgrade actions. These times vary depending on your system and the particular upgrade you are performing. Schedule an appropriate service interval time with the customer.

Table 1. Estimated Time to Upgrade Management Nodes and Diskless Nodes

| Upgrade Action | Estimated Time |
|--|----------------|
| Upgrade Primary Management | 26 minutes |
| Upgrade Secondary Management | 14 minutes |
| Upgrade First Diskless Node | 15 minutes |
| Upgrade Remaining Diskless Nodes (including CNG) | 14 minutes |

Table 2. Estimated Time to Upgrade Software

| Upgrade Action | Estimated Time |
|--|----------------|
| Upgrade Sonexion Software – Preparation Time (System Up) | 4 hours |
| Upgrade Sonexion Software – Upgrade Time (System Down) | 1.5 hours* |
| Apply 1.5.0 System Updates | As Needed |

* This estimate applies to a Sonexion 1600 system with 2 SSUs. Add approximately 1 minute for each additional 3 nodes

Table 3. Estimated Time to Upgrade Firmware (Post Upgrade)

| Upgrade Action | Estimated Time |
|---|--|
| Leveling of the USM Firmware on the MMU and SSUs | Depends on the size of the cluster: Single SSU - about 40 minutes. Each additional SSU - add 20 minutes. |
| Leveling of the LSI Firmware | 1 hour |
| Non-Live Leveling of the Mellanox HCA Firmware MMU, SSUs, and CNG | 1 hour |
| For Sonexion 1600, Leveling the Mellanox InfiniBand Switch Firmware (6025, 5023, or 5035) | 1 hour (2 switches) |
| Leveling the Firmware Running on the HDDs and SSDs | 1 hour |

Sonexion 900 systems are shipped to customers without network switches (IB or 40GbE). Customers must provide their own. Therefore, the time estimate above for leveling Mellanox InfiniBand switch firmware applies to the Sonexion 1600 only. If a Sonexion 900 customer does use the same Mellanox InfiniBand switches that are qualified by Cray, it is recommended that the switch firmware be leveled to the Cray-qualified firmware level. If this is done, the estimated time is the same as for the Sonexion 1600.

Upgrade Scope and Limitations

Scope:

- The only supported upgrade path covered by this procedure is from software release 1.4.0-28, plus any level of 1.4.0 System Update (SU), to 1.5.0-38.
- It is not backward compatible.

Table 4. Supported Upgrade Path

| Software Version-Build |
|------------------------|
| 1.4.0-28 to 1.5.0-38 |

Limitations:

-
- Software release 1.4.0-28 must be installed to enable the 1.5.0 upgrade.
- The upgrade procedure is packaged as a script that runs a sequence of CLI commands. There is no GUI option available at this time.
- The upgrade script will upgrade OS-level components. Firmware upgrades are applied separately, as described in this document.
- There is no supported downgrade path.

Post-Installation Steps after Upgrading Software to 1.5.0

The following post-installation steps must be performed after the software upgrade script has completed successfully:

- Install any available Sonexion 1.5.0 system updates.
- Level the Mellanox HCA, LSI, USM, and HDD/SSD firmware.
- Level the Mellanox InfiniBand switch firmware.

The above steps are described in [Post-Installation Steps](#) on page 28.

Reference Documents

The reference documents listed in this section may be downloaded from CrayPort or obtained from Cray Support.

Leveling of the USM Firmware on the MMU and SSUs

These procedures update Sonexion 1600/900 1.5.0 systems with new firmware for all Cray enclosures and modules. See *Sonexion 1600 USM Firmware Update Guide 1.5* and *Sonexion 900 USM Firmware Update Guide 1.5*.

Leveling of the LSI Firmware

This procedure updates Sonexion 1600/900 1.5.0 systems with new firmware for all LSI SAS controllers in the system. Contact Cray Support to obtain the appropriate proprietary manual for upgrading LSI firmware.

Non-Live Leveling of the Mellanox HCA Firmware on the MMU and SSUs

This procedure checks the Mellanox firmware version on the 2U quad server HCAs (in the MMU and CNG) and OSS HCAs (in each SSU) for Sonexion 1600/900 1.5.0 systems, and, if necessary, updates the firmware to the required level. Contact Cray Support to obtain the appropriate proprietary manual for upgrading Mellanox HCA firmware.

Leveling the Mellanox InfiniBand Switch Firmware

Documents for updating firmware on Mellanox network switches (InfiniBand) can be obtained from the same source used for this publication. See *Sonexion 1600 Mellanox InfiniBand Switch Firmware Update 1.5* and *See Sonexion 900 Mellanox InfiniBand Switch Firmware Update 1.5*.

Some Sonexion 900 systems do not ship from Cray with high-speed network switches. Instead, the switches are provided by the customer. If the customer uses the same model network switches that are qualified by Cray for use in Sonexion systems, it is recommended that the firmware be at the same level.

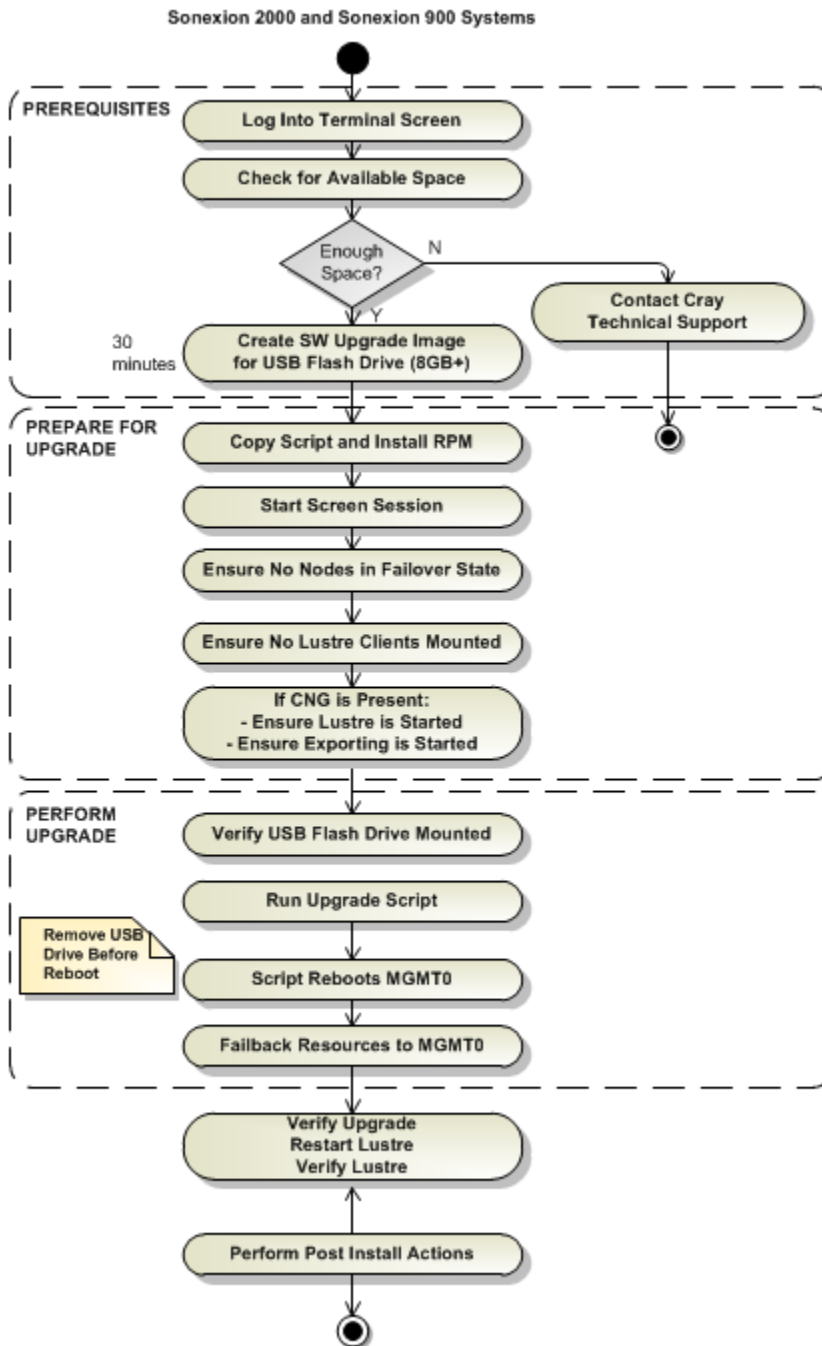
Leveling the Firmware Running on hard disks and SSDs

For more information, contact Cray Support to obtain the appropriate proprietary manual for upgrading firmware on hard disks and SSDs.

Process Overview

The following flowchart provides a visual overview of the software upgrade process.

Figure 1. Sonexion 1.5 Upgrade Flowchart



Prerequisites

The following conditions must be met for a successful system upgrade:

- Must have the latest System Update (SU) for 1.5.0 available to pre-load on the cluster during the preparation steps (see [Prepare the Cluster with the Latest System Update for 1.5.0](#)) before running the upgrade process.
- System Update 1.5.0-007 or greater should be used. To obtain System Updates for Release 1.5.0 and Release Notes with instructions, contact Cray Support.
- Must have available free space of at least 6GB at the root file system and 8GB on shared storage.
- Must have the following required files. (See [Create the Software Upgrade Image for the USB Flash Drive](#) and [Create the Software Upgrade Image from Loopback Disk Image](#).)
 - Engineering Change bundle:

CS_1.5.0-38-x86_64-USB.img
 - (Optional, if the ISO-based upgrade will be performed) ISO-based upgrade DVD image:

CS_1.5.0-38-x86_64-DVD.iso
 - Latest upgrade script package: The current filename as of the release date of this document is:

cs-upgrade-1.0.x.1.5-36.d25925fb.noarch.rpm
- All commands in the software upgrade procedure require root privileges.
- The software upgrade procedure assumes an error-free 1.4.0 installation.
- Ensure no one is logged into the system at the time of the software upgrade.
- Allocate a free network to be used for the secondary IPMI interfaces. This network must not overlap with any of the directly attached networks of the cluster. To find directly attached networks check output from the command:

```
[root@n000]# ip a l | awk '$1 == "inet" {print $2}' | sort
```

In this sample output, none of the displayed networks may be used as a secondary IPMI network:

```
10.10.0.2/16
10.10.0.3/16
10.106.54.154/22
127.0.0.1/8
172.16.2.1/16
172.16.2.3/16
172.18.1.18/24
```



WARNING: Do not start the upgrade procedure until all the prerequisites have been met. Use the checklist in the checklist below, to verify the prerequisites have been met.

Table 5. Prerequisite Checklist

| Prerequisite | Section | Complete Yes/No |
|--|---|-----------------|
| <i>Preparation for Upgrade:</i> | | |
| Verify that current software version is supported for the upgrade. | Upgrade Scope and Limitations | |
| Verify that the estimated interval to perform the system upgrade is available. | Estimated Service Interval | |
| Verify Cray Support is aware of upgrade execution date and time. | Open a Ticket with Cray Technical Support | |
| Verify that there is enough available free space in the root file system and shared storage. | Prepare the Cluster with the Latest System Update for 1.5.0 | |
| Verify required files are in place. | Prerequisites (see bullet list): Create the Software Upgrade Image for the USB Flash Drive | |
| Verify software upgrade image is useable and mounted on the primary MGMT node. | Create the Software Upgrade Image for the USB Flash Drive | |
| <i>Before Running the Upgrade Script:</i> | | |
| Install ISLs between management switches | Install Inter-switch Links (ISL) | |
| Configure Spanning Tree and DHCP on management switches | Configure Spanning Tree and DHCP on Brocade Management Switches | |
| Prepare to run the script. | Prepare to Run the Upgrade Script | |
| Verify system is healthy and no nodes have failed over. | Run Failback before Starting the Upgrade | |

Begin Upgrade

Requirements:

- Use the Linux screen TTY multiplexor to protect from a possible disconnect problem. If the connection is lost while the software upgrade script runs, use the screen commands described below to reconnect to the session.
- Open an additional terminal for access to a separate shell as needed by prompts in the software upgrade script.

Subtasks:

- [Prepare the Cluster with the Latest System Update for 1.5.0](#)
 - [Check for Available Space](#)
 - [Create the Software Upgrade Image for the USB Flash Drive](#)
 - [Create the Software Upgrade Image from Loopback Disk Image](#)
1. From `putty.exe` or another terminal program that can be used to access the cluster, open a terminal shell and log in to the primary MGMT node:

```
[CLIENT]$ ssh admin@primary_MGMT_node
```

2. Change to root user:

```
[admin@n000]$ sudo su -
```

3. Open a screen session:

```
[root@n000]# screen [-S screen_session_name] -L
```

4. To disconnect, press **Ctrl-a**, then **d**
5. To obtain the ID of the session that was disconnected, type:

```
[root@n000]# screen -list
```

Here is an example command and output:

```
[root@snx11000n005-oem ~]# screen -list
There are screens on:
10456.pts-10.snx11000n5-oem (Attached)
12450.pts-0.snx11000n5-oem (Detached)
2 Sockets in /var/run/screen/S-root.
```

6. Reconnect to the terminal session:

```
[root@n000]# screen -rd id_from_the_screen-list
```

From this point forward, if the connection is interrupted while the software upgrade script is running, the terminal session at the cluster should still be available. Perform Steps 5 and 6 again to reconnect.

Prepare the Cluster with the Latest System Update for 1.5.0

To ensure proper upgrade support for clusters that may have any 1.4.0 System Update (SU) installed, the latest 1.5.0 SU should be prepared on the cluster before performing the software upgrade from 1.4.0 to 1.5.0. The latest 1.5.0 SU may be obtained from CrayPort.

The latest System Update for 1.5.0 needs to be applied (not simply "prepared") **after** the upgrade script has completed. See [Apply 1.5.0 System Updates If Applicable](#) on page 28 for more information.

The following steps are taken from the 1.5.0 System Update Release Notes.

7. If not already in an SSH session with the primary MGMT node, connect to the primary MGMT node via SSH:

```
[Client]$ ssh -l admin MGMT0
```

8. Change to root user:

```
[MGMT0]$ sudo su -
```

9. Identify the MGMT node hosting NFS services. It is the node hostname next to the "Started" status:

```
[MGMT0]# crm_mon -l | grep nfsd
```

Sample output:

```
[root@snx11000n000 ~]# crm_mon -l | grep nfsd
prn-nfsd (ocf::heartbeat:nfsserver): Started snx11000n001
```

The node indicated as started is likely MGMT1, but the specific MGMT node on which this service is running is unimportant. The remaining instructions assume MGMT1.

10. Log into the MGMT NFS server node by its hostname via SSH:

```
[MGMT0]# ssh hostname_from_previous_step
```

Sample output:

```
[root@snx11000n000 ~]# ssh snx11000n001
Last login: Sat May 31 15:43:26 2014 from 172.16.0.1
```

11. Start a screen session. A log will be kept in /root/screenlog.0:

```
[MGMT1]# screen -L -S update
```

If the SSH session is interrupted, the update script should continue to run in the screen session. To reconnect, do the following:

- a. Log back in to the NFS host node via SSH.
- b. Run `screen -ls` to confirm the screen is still active.
- c. To reconnect to the screen session, run:

```
[MGMT1]# screen -L -RDS update
```

12. Copy `system-update_cs-1.5.0-001.12.sh` from XIC to `/tmp` on the MGMT1 node.

13. Make the script executable:

```
[MGMT1]# chmod +x /tmp/system-update_cs-1.5.0-001.12.sh
```

14. Prepare the system to install the bundle as follows. This differs from the System Update Release Notes by including the `--force` option.

```
[root@n001]# /tmp/system-update_cs-1.5.0-007.71.sh prepare --force
```

At this point, the cluster has been prepared for proper upgrade support. Stop after running the prepare step above. **Do not** perform any additional steps from the full System Update process. Proceed with the remaining steps to upgrade Sonexion software from 1.4.0 to 1.5.0.

Check for Available Space

The software upgrade procedure assumes an error-free 1.4.0 installation.

The upgrade procedure requires a minimum of 6GB of free space available on the root file system level and a minimum of 8GB on shared storage. The upgrade script will automatically check to ensure that there is enough available space to perform the upgrade. If there is not enough space the upgrade script will print an error message and stop. You will then need to clean space and restart the upgrade script.

Use the following procedure to manually check available space before you actually run the upgrade script.

15. Check available space on the root file system:

```
[root@n000]# df -h /
```

If there is less than 6 GB on the root file system, you do not have enough space. STOP. Contact Cray Technical Support.

16. Check available space on the shared storage:

```
[root@n000]# ssh nfsserv df -h /mnt/nfsdata
```

If there is less than 8 GB on the shared storage, you do not have enough space. STOP. Contact Cray Technical Support.

17. Create the Software Upgrade Image, using one of the two methods for creating the software image:

- USB (recommended). See [Create the Software Upgrade Image for the USB Flash Drive](#), below.
- Using a loopback disk image. See [Create the Software Upgrade Image from Loopback Disk Image](#).

Create the Software Upgrade Image for the USB Flash Drive

The software upgrade image is a USB image file (.img) contained in an EC bundle, which is not a format the Sonexion 1600/900 system can access. To make the image usable for the system, you must convert it using the Linux dd command and then install it on a USB flash drive.

The USB flash drive must have a capacity of at least 8GB.

In the example commands and sample output in this section, the primary MGMT node is "snx11000n000". When you run these commands, replace this value with the appropriate hostname or IP address.

- 18.** Locate and download the USB software image of the Sonexion 1600/900 1.5.0 EC bundle. Also, download the latest version of the upgrade script package.

- CS_1.5.0-38-x86_64-USB.img
- cs-upgrade-1.0.x.1.5-36.d25925fb.noarch.rpm (This is the current upgrade script package name as of the release date of this document.)

The EC bundle is managed through the EC process. Contact Cray Support to obtain the correct image file.

- a. Download the files to your workstation or terminal program's home directory.
- b. Use **scp** to copy the files to the /mnt/mgmt directory on the primary MGMT node:

```
[Client]$ scp CS_1.5.0-38-x86_64-USB.img admin@snx11000n000:/mnt/mgmt
[Client]$ scp cs-upgrade-1.0.x.1.5-36.d25925fb.noarch.rpm
admin@snx11000n000:/mnt/mgmt
```

For a Windows client, optionally download and use WinSCP to copy the files. WinSCP is an open source SFTP client, SCP client, FTPS client, and FTP client for Windows.

- c. Log in to the primary MGMT node:

```
[Client]$ ssh admin@snx11000n000
```

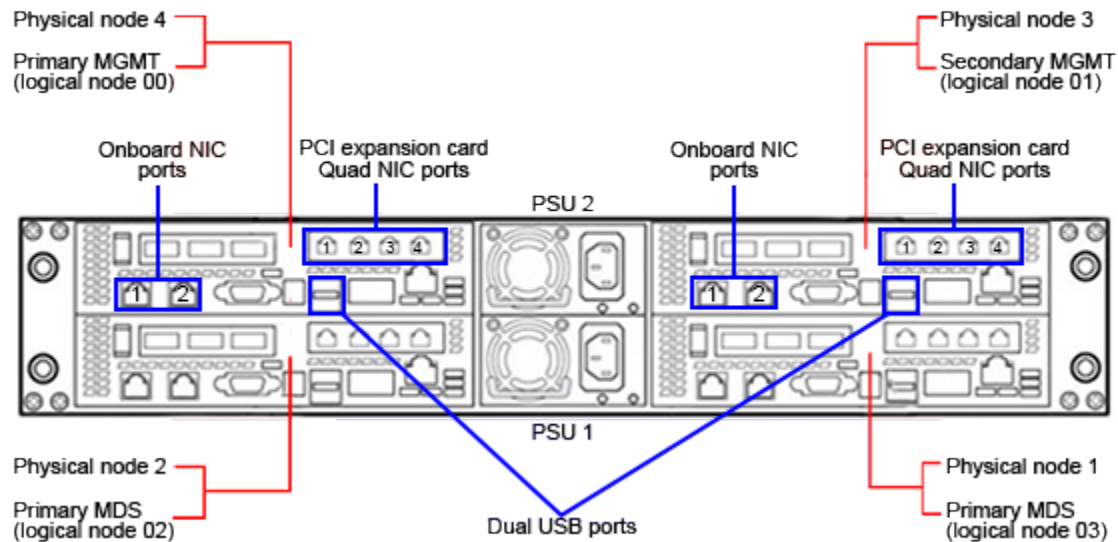
- d. Change to root user:

```
[snx11000n000]$ sudo su -
```

- 19.** Insert the USB flash drive:


- For the Sonexion 1600, insert the flash drive into a USB port at the rear of the server hosting the primary MGMT node.
- As shown in the following figure, there are dual USB ports on the rear panel of each server in the 2U quad server.

Figure 2. Sonexion 1600 Quad Server Ports



- For the Sonexion 900, insert the USB flash drive into a USB port at the rear of the Sonexion 900 OSS hosting the primary MGMT node. This is the upper OSS in the MMU. See the following figure for reference.

Figure 3. OSS Rear Panel on Sonexion 900



- Find the correct device name of the USB flash drive by running the `lsusb` command on the primary MGMT node (where the USB flash drive was plugged in).
- Transfer the software image to the USB flash drive using a low-level tool, such as `dd`, that can transfer raw blocks. For example, run:

```
[snx11000n00]# dd if=/mnt/mgmt/CS_1.5.0-38-x86_64-USB.img of=<USB device name> oflag=direct bs=1M
```

Transferring the software image will take about 25 to 30 minutes. This is a sample `dd` command and output:

```
[root@snx11000n000 ~]# dd if=/mnt/mgmt/CS_1.5.0-38-x86_64-USB.img of=/dev/sdc oflag=direct bs=1M
3745+0 records in
```

```
3745+0 records out
3926917120 bytes (3.9 GB) copied, 9.44628 s, 416 MB/s
```

22. Remove and then reinsert the USB flash drive into the USB port of the primary MGMT node. This step is required to ensure correct mounting of the USB flash drive before performing the software upgrade.

23. Verify that the USB flash drive is mounted:

```
[root@n000]# mount | grep media
```

Here is an example command and output:

```
[root@snx11000n000 ~]# mount | grep media
/dev/sdaa1 on /media/SLX_38_USB type ext3
(rw,relatime,errors=continue,user_xattr,acl,barrier=1,data=ordered)
/dev/loop0 on /mnt/media/updates type iso9660 (ro,relatime)
```

24. Skip to Upgrade Sonexion 1600/900 Software

Create the Software Upgrade Image from Loopback Disk Image

25. Locate and download the ISO-based upgrade DVD image of the Sonexion 1600/900 1.5.0 EC bundle.

- CS_1.5.0-38-x86_64-DVD.iso
- cs-upgrade-1.0.x.1.5-36.d25925fb.noarch.rpm (This is the current upgrade script package name as of the release date of this document.)

The EC bundle is managed through the EC process. Cray FAEs and Support personnel can download the file from CrayPort. It is recommended that customers coordinate their software upgrade with Cray Support.

26. Download the files to the workstation or terminal program's home directory.

27. Use scp to copy the files to the /mnt/mgmt directory on the primary MGMT node:

```
[Client]$ scp CS_1.5.0-38-x86_64-DVD.iso snx11000n000:/mnt/mgmt
[Client]$ scp cs-upgrade-1.0.x.1.5-36.d25925fb.noarch.rpm admin@snx11000n000:/mnt/mgmt
```

For a Windows client, optionally download and use WinSCP to copy the files. WinSCP is an open source SFTP client, SCP client, FTPS client, and FTP client for Windows.

28. Log in to the primary MGMT node:

```
[Client]$ ssh admin@snx11000n000
```

29. Change to root user:

```
[snx11000n000]$ sudo su -
```

30. Change to the /mnt/mgmt directory:

```
[snx11000n000]# cd /mnt/mgmt
```

Mount the loopback device:

```
[snx11000n000]# mkdir -p /mnt/media/updates  
[snx11000n000]# mount -o ro,loop /mnt/mgmt/CS_1.5.0-38-x86_64-DVD.iso /mnt/  
media/updates
```

Upgrade Sonexion 1600/900 Software



WARNING: Before starting the upgrade procedure, note the following:

- Ensure that the MGMT stack is running on the primary MGMT node (`crm_mon -1rQ`, etc.).
- Make certain no other users are logged into the Sonexion 1600/900 system when the upgrade procedure is performed. If the system includes CNG nodes, Lustre must be running.
- Make sure that exporting is started if the system contains CNG nodes.
- Perform the software upgrade procedure on the primary MGMT node (`snx11000n000`).

All commands in the software upgrade procedure require root privileges.

Subtasks:

- [Install Inter-switch Links \(ISL\)](#)
- [Configure Spanning Tree and DHCP on Brocade Management Switches](#)
- [Prepare to Run the Upgrade Script](#)
- [Run the Software Upgrade Script](#)
- [Script Actions](#)
- [Steps Following Script](#)
- [Verify and Confirm the Software Upgrade](#)

Install Inter-switch Links (ISL)

The two management switches in the base rack and all expansion racks must be connected to each other with an inter-switch link (ISL) cable before running the upgrade script.

1. Determine if ISL cables are already installed between the two management switches in the base rack and in each of the expansion racks.

If the cables are already installed, skip to [Configure Spanning Tree and DHCP on Brocade Management Switches](#), below.

2. Install the ISL cables between the two management switches in the base rack and in each of the expansion racks.

The recommended port assignments for the ISL on the management switches is: For 24-port switches, connect port 19 on SW0 to port 19 on SW1.

For 48-port switches, connect port 41 on SW0 to port 41 on SW1.

For some configurations, there may be an existing cable in the recommended ISL port. Refer to the "Internal Cabling" section of the *Sonexion 2000 Field Installation Guide* or *Sonexion 900 Field Installation Guide*, and relocate the existing cable to its proper place. The ISL can then be installed in the recommended port. Do **not**

remove the ISL cables when the upgrade is complete. The ISL cables must remain in place since they are required for proper functioning of the Dual LMN feature in Sonexion 1.5.0.

Configure Spanning Tree and DHCP on Brocade Management Switches

If the Sonexion system uses Brocade ICX 6430 management switches, configure all of them to use Spanning Tree Protocol and DHCP. This section describes how to perform this configuration, which must be done before running the upgrade script.

No special configuration is required if the Sonexion system uses NetGear management switches. Spanning Tree Protocol and DHCP are enabled by default.

Required Equipment: Verify that the following equipment is available before configuring the management switches.

- Console with monitor and keyboard OR
PC/Laptop with a DB9 serial COM port, mouse, keyboard. The serial port should be configured for 9600 Kbps, 8 data bits, no parity and 1 stop bit.
- Rollover/Null modem serial cable (DB9 to RJ45 Ethernet)
- Serial terminal emulator program, such as SecureCRT, PuTTY, Tera Term, Terminal Window (Mac) etc.
- ESD strap, boots, garment or other approved methods

Perform the following procedure for each Brocade management switch in the system.

3. To connect to the switch, connect the rollover cable to the console port (RJ45) at the left side of the Management Switch and to the serial COM port (DB9) on the PC/Laptop. Refer to the following figure.

Figure 4. Console Port on Brocade ICX 6430 Switch



4. To use a PuTTY/SecureCRT emulator program, open a terminal session and specify the following serial port settings:
Baud rate (bits per second): 9600
Data bits: 9600
Stop bits: 1
Parity: None
Flow control: Off
If using a terminal window on a Macintosh, type the following and press **Enter**:

```
screen /dev/ttyUSB0 cs8,9600
```

After connecting to the switch, the terminal screen should display the switch command prompt, which is shown as [switch] in the remaining commands.

5. Enter the following commands at the switch prompt to perform a minimal configuration. After the software upgrade is completed, a switch configuration script will be run to fully configure the switch. The commands below illustrate configuration for a 24-port Brocade management switch. 48 port switches will have the same configuration except for port 41 (the ISL port) and/or any uplink ports that connect to expansion racks.

```
[switch]enable
[switch]configure terminal
[switch]spanning-tree 802-1w enable
[switch]interface Ethernet 1/1/19
```

Port number 19 is used when installing the ISL cables between 24-port management switches. This is 41 when configuring a 48-port Brocade management switch. Refer to [Install Inter-switch Links \(ISL\)](#), above.

```
[switch]interface Ethernet 1/1/1 to 1/1/18
[switch]spanning-tree 802-1w admin-edge-port
[switch]exit
[switch]interface Ethernet 1/1/20 to 1/1/24
[switch]spanning-tree 802-1w admin-edge-port
[switch]exit
```

Use the following three additional commands only for configuring 48-port switches.

```
[switch]interface Ethernet 1/1/25 to 1/1/48
[switch]spanning-tree 802-1w admin-edge-port
[switch]exit
```

```
[switch]ip dhcp-client enable
[switch]ip dhcp-client auto-update enable
[switch]write memory
[switch]show ip address
```

Make note of the IP address that is displayed, to use for telnet access later.

6. Disconnect the PC or terminal from the switch.
7. Repeat Steps 3 to 6 for the remaining Brocade management switches.

Prepare to Run the Upgrade Script

8. Confirm that the upgrade script package was copied to `/mnt/mgmt` on the primary MGMT node. If it was not copied, see step [18.b](#) on page 13 of [Begin Upgrade](#) on page 10, to copy the file.
9. Install the upgrade script package. On the primary MGMT node, run:

```
[root@n000]# yum install -y /mnt/mgmt/cs-upgrade-1.0.x.
1.5-36.d25925fb.noarch.rpm
```

10. Start the Upgrade Script screen session. Using the terminal screen is required to run the upgrade 1.5.0 script. Start a screen session if it is not already running:

```
[root@n000]# screen -L -S user_specified_session_name
```

11. In the following steps, determine if any nodes are in a failed-over state. If so, fail back the nodes before starting the upgrade process.

Run `crm_mon -l` on the primary MGMT node to verify the resource state.

```
[root@n000]# crm_mon -l
```

Following is an example output for a Sonexion 1600 system:

```
[root@snx11000n000 ~]# crm_mon -l
=====
Last updated: Sun Mar 29 19:52:19 2015
Last change: Sun Mar 29 19:41:36 2015 via crm_resource on snx11000n001
Stack: Heartbeat
Current DC: snx11000n001 (b56c0186-d299-4254-9f55-1e56b6b6e8f7) - partition
with quorum
Version: 1.1.6.1-6.el6-0c7312c689715e096b716419e2ebc12b57962052
2 Nodes configured, unknown expected votes
40 Resources configured.
=====
Online: [ snx11000n001 snx11000n000 ]
  baton (ocf::heartbeat:baton): Started snx11000n001
  snx11000n000_mdadm_conf_regenerate (ocf::heartbeat:mdadm_conf_regenerate):
Started snx11000n000
  snx11000n001_mdadm_conf_regenerate (ocf::heartbeat:mdadm_conf_regenerate):
Started snx11000n001
  snx11000n000-stonith (stonith:external/ipmi): Started snx11000n000
  snx11000n001-stonith (stonith:external/ipmi): Started snx11000n001
  prm-httpd (lsb:httpd): Started snx11000n000
  prm-mysqld (lsb:mysqld): Started snx11000n000
  prm-nfslock (lsb:nfslock): Started snx11000n001
  prm-bebundd (lsb:bebundd): Started snx11000n000
  Clone Set: cln-cerebrod [prm-cerebrod]
    Started: [ snx11000n001 snx11000n000 ]
  prm-conman (lsb:conman): Started snx11000n000
  prm-dhcpd (lsb:dhcpd): Started snx11000n001
  prm-xinetd (lsb:xinetd): Started snx11000n001
  Clone Set: cln-syslogng [prm-syslogng]
    Started: [ snx11000n001 snx11000n000 ]
  prm-nodes-monitor (lsb:nodes-monitor): Started snx11000n000
  Clone Set: cln-ses_mon [prm-ses_monitor]
    Started: [ snx11000n001 snx11000n000 ]
  Clone Set: cln-nsca_passive_checks [prm-nsca_passive_checks]
    Started: [ snx11000n001 snx11000n000 ]
  Resource Group: grp-icinga
    prm-icinga (lsb:icinga): Started snx11000n000
    prm-nsca (lsb:nsca): Started snx11000n000
    prm-npcd (lsb:npcd): Started snx11000n000
  Resource Group: grp-plex
    prm-rabbitmq (lsb:rabbitmq-server): Started snx11000n000
    prm-plex (lsb:plex): Started snx11000n000
  prm-repo-local (ocf::heartbeat:Filesystem): Started snx11000n001
  prm-repo-remote (ocf::heartbeat:Filesystem): Started snx11000n000
  prm-db2puppet (ocf::heartbeat:oneshot): Started snx11000n000
  Clone Set: cln-puppet [prm-puppet]
    Started: [ snx11000n001 snx11000n000 ]
  prm-nfsd (ocf::heartbeat:nfsserver): Started snx11000n001
  prm-vip-eth0-mgmt (ocf::heartbeat:IPAddr2): Started snx11000n000
  prm-vip-eth0-nfs (ocf::heartbeat:IPAddr2): Started snx11000n001
  Resource Group: snx11000n000_md64-group
    snx11000n000_md64-raid (ocf::heartbeat:XYRAID): Started snx11000n000
```

```

snx11000n000_md64-fsys (ocf::heartbeat:XYMNTR): Started snx11000n000
snx11000n000_md64-stop (ocf::heartbeat:XYSTOP): Started snx11000n000
Resource Group: snx11000n000_md67-group
snx11000n000_md67-raid (ocf::heartbeat:XYRAID): Started snx11000n001
snx11000n000_md67-fsys (ocf::heartbeat:XYMNTR): Started snx11000n001
snx11000n000_md67-stop (ocf::heartbeat:XYSTOP): Started snx11000n001

```

Following is sample output for a Sonexion 900 system:

```

[root@snx11000n000 ~]# crm_mon -l
=====
Last updated: Tue Mar 24 08:32:10 2015
Last change: Tue Mar 24 07:41:00 2015 via cibadmin on snx11000n001
Stack: Heartbeat
Current DC: snx11000n001 (61d820a9-dffd-480b-a446-cff5d66bed03) - partition
with quorum
Version: 1.1.6.1-6.el6-0c7312c689715e096b716419e2ebc12b57962052
2 Nodes configured, unknown expected votes
46 Resources configured.
=====
Online: [ snx11000n000 snx11000n001 ]
Clone Set: clone-gem-ipmi-stonith [gem-ipmi-stonith]
Started: [ snx11000n000 snx11000n001 ]
snx11000n001_mdadm_conf_regenerate (ocf::heartbeat:mdadm_conf_regenerate):
Started snx11000n001
snx11000n000_mdadm_conf_regenerate (ocf::heartbeat:mdadm_conf_regenerate):
Started snx11000n000
baton (ocf::heartbeat:baton): Started snx11000n001
snx11000n001_ibstat (ocf::heartbeat:ibstat): Started snx11000n001
snx11000n000_ibstat (ocf::heartbeat:ibstat): Started snx11000n000
Resource Group: snx11000n001_md66-group
snx11000n001_md0-raid (ocf::heartbeat:XYRAID): Started snx11000n001
snx11000n001_md66-raid (ocf::heartbeat:XYRAID): Started snx11000n001
snx11000n001_md66-fsys (ocf::heartbeat:XYMNTR): Started snx11000n001
snx11000n001_md66-stop (ocf::heartbeat:XYSTOP): Started snx11000n001
prm-httpd (lsb:httpd): Started snx11000n000
prm-mysqld (lsb:mysqld): Started snx11000n000
prm-nfslock (lsb:nfslock): Started snx11000n001
prm-bebundd (lsb:bebundd): Started snx11000n000
Clone Set: cln-cerebrod [prm-cerebrod]
Started: [ snx11000n001 snx11000n000 ]
prm-conman (lsb:conman): Started snx11000n000
prm-dhcpd (lsb:dhcpd): Started snx11000n001
prm-xinetd (lsb:xinetd): Started snx11000n001
Clone Set: cln-syslogng [prm-syslogng]
Started: [ snx11000n000 snx11000n001 ]
prm-nodes-monitor (lsb:nodes-monitor): Started snx11000n000
Clone Set: cln-ses_mon [prm-ses_monitor]
Started: [ snx11000n000 snx11000n001 ]
Clone Set: cln-nsca_passive_checks [prm-nsca_passive_checks]
Started: [ snx11000n000 snx11000n001 ]
Resource Group: grp-icinga
prm-icinga (lsb:icinga): Started snx11000n000
prm-nsca (lsb:nsca): Started snx11000n000
prm-npcd (lsb:npcd): Started snx11000n000
Resource Group: grp-plex
prm-rabbitmq (lsb:rabbitmq-server): Started snx11000n000
prm-plex (lsb:plex): Started snx11000n000
prm-repo-local (ocf::heartbeat:Filesystem): Started snx11000n001
prm-repo-remote (ocf::heartbeat:Filesystem): Started snx11000n000

```

```

prm-db2puppet (ocf::heartbeat:oneshot): Started snx11000n000
Clone Set: cln-puppet [prm-puppet]
Started: [ snx11000n000 snx11000n001 ]
prm-nfsd (ocf::heartbeat:nfsserver): Started snx11000n001
prm-vip-eth0-mgmt (ocf::heartbeat:IPaddr2): Started snx11000n000
prm-vip-eth0-nfs (ocf::heartbeat:IPaddr2): Started snx11000n001
Resource Group: snx11000n000_md64-group
  snx11000n000_md64-raid (ocf::heartbeat:XYRAID): Started snx11000n000
  snx11000n000_md64-fsys (ocf::heartbeat:XYMNTR): Started snx11000n000
  snx11000n000_md64-stop (ocf::heartbeat:XYSTOP): Started snx11000n000
Resource Group: snx11000n000_md67-group
  snx11000n000_md67-raid (ocf::heartbeat:XYRAID): Started snx11000n001
  snx11000n000_md67-fsys (ocf::heartbeat:XYMNTR): Started snx11000n001
  snx11000n000_md67-stop (ocf::heartbeat:XYSTOP): Started snx11000n001

```

12. To rebalance the HA resources because they were failed over, connect and run:

```
[root@n000]# ssh puppet '/opt/xyratex/bin/cscli failback -n primary MGMT node'
```

Here is an example command:

```
[root@n000]# ssh puppet '/opt/xyratex/bin/cscli failback -n snx11000n000'
```

13. Unmount all Lustre clients from the Lustre file system.

There is no need to stop Lustre on the nodes, although it should not be mounted on clients. Lustre will be stopped before the upgrade script completes. If the cluster contains CNG nodes, Lustre and exporting must be started.

14. Start exporting for CNG nodes:

```

[snx11000n000]# cscli mount -f fsname
[snx11000n000]# cscli export enable -y
[snx11000n000]# cscli export apply -y

```

Run the Software Upgrade Script

NOTE:

- Reminder: As noted earlier, the upgrade procedure assumes an error-free 1.4.0 installation.
- Removing the USB flash drive: During the procedure, manually remove the USB flash drive **before** the primary MGMT node reboots. See step 18 on page 24 for the proper point to do this. Do not leave the upgrade procedure running unattended until the USB flash drive is removed.

15. Start the upgrade script and ensure that it can accurately read the ports:

```
[root@n000]# /opt/xyratex/upgrade/upgrade.sh -a -b
```

Brief sample output:

```

# /opt/xyratex/upgrade/upgrade.sh -a -b
upgrade INF: Exporting persistent CTDB databases
Please enter network for the secondary IPMI addresses (default is
10.10.0.0/16):

```

This statement checks the system to ensure there is enough storage space and that the USB is plugged in.

NOTE: On the Sonexion 900:

- After the Sonexion 1.5 prepare command is run, the SU files may need to be removed, both zipped and unzipped, to free enough space on the system.
- If node 0 is bad, error checking automatically reports node1 and no USB. If node 1 is bad it shows node 0 as a good node.

16. Enter the network allocated for the secondary IPMI interfaces or press **Enter** (or **Return**) to use default one (10.10.0.0/16).

```
10.10.0.0/16 is going to be used as the secondary IPMI network. Please hit
Return to confirm or type 'quit' to quit now
```

17. Press **Enter** (or **Return**) to confirm the selected network or type quit to terminate the upgrade process. If desired, it is safe to quit now and restart the upgrade later.

```
Total size of LMT data:
2.5M    /mnt/mgmt/var/lib/mysql/filesystem_testfs
Type: 'recreate' or 'repair' to choose LMT handling method, 'quit' to
terminate upgrade process now, or 'help' to learn about available methods,
and press 'Enter':
```

Script Actions

The script will repair or recreate the LMT database later, but at this point one of the following methods must be selected. In either case, a complete backup of all databases will be created.

recreate: Unconditionally drops the database and recreates it from scratch. This is the recommended method if historical Lustre performance metrics are not needed, or if the LMT size is large enough to affect system downtime because of the length of the check/repair process.

repair: Runs **mysqlcheck** in repair mode. This method preserves historical Lustre performance metrics but may take a long time to run, depending on the actual LMT database size. If desired, it is also safe to quit now and restart the upgrade later.

From this point forward, the upgrade is automatic. Samples of expected output are shown below, but do not require interaction to proceed. Removing the USB flash drive once it is done being used is the only action needed. Following is an example upgrade output.

NOTE: Watch for the status statement "Updates were transferred successfully...", which is underscored below. The script does not stop to make a notification of this.

```
prm-nfsd      (ocf::heartbeat:nfsserver) Started
...
  eval1000_md67-fsys (ocf::heartbeat:XYMNTR) Started
  eval1000_md67-stop (ocf::heartbeat:XYSTOP) Started
---Dots below this line indicate close finish ( 12)---
.....
upgrade INF: New release packages will be copied now onto the shared
storage. This may take 3-10 minutes, depending on USB drive speed
upgrade NXT: prepare_update: /opt/xyratex/bin/cscli prepare_update --run
Updates were transferred successfully to the install location. Your USB
drive may now be safely removed.
Old version: 1.4.0-28
New version: 1.5.0-38
```

```

upgrade INF: Refreshing yum repository metadata
upgrade NXT: prepare_update: pdsh -g mgmt yum clean metadata
eval1000: Cleaning repos: cs updates_base updates_firmware
eval1000: 0 metadata files removed
eval1000: 0 sqlite files removed
eval1000: 0 metadata files removed
eval1001: Cleaning repos: cs updates_base updates_firmware
eval1001: 9 metadata files removed
eval1001: 3 sqlite files removed
eval1001: 0 metadata files removed
upgrade NXT: prepare_update: createrepo /mnt/nfsdata/release/cs-1.4.0-28

```

Steps Following Script

18. When `/opt/xyratex/bin/cscli prepare_update` finishes, remove the USB drive from the primary MGMT node to avoid booting from it on reboot. It is safe to remove the drive any time from this point, but it must be removed before the primary MGMT node reboots.

```

...
upgrade INF: Rebooting primary mgmt, eval1000. After the node boots, please
reconnect and run 'ssh eval1001 /opt/xyratex/bin/cscli failback -n eval1000'
to rebalance HA resources
upgrade NXT: reboot_mgmt_prim: reboot
Broadcast message from root@eval1000
(/dev/pts/7) at 13:11 ...
The system is going down for reboot NOW!
Timeout, server not responding.

```

The script reboots the primary MGMT node, the screen session terminates, and connectivity is lost. Remove the USB drive before the primary MGMT node reboots to avoid booting from it upon reboot. If the ISO was mounted via loopback, it is automatically unmounted and no further action is necessary.

19. Once the primary MGMT node is back online, log in to it again via SSH:

```
[CLIENT]$ ssh -l admin_primary_MGMT_node
```

20. Change to root user:

```
[admin@n000]$ sudo su -
```

21. Fail back the primary MGMT node's resources from the secondary MGMT node. On the primary MGMT node, run:

```
[root@n000]# ssh puppet '/opt/xyratex/bin/cscli failback -n MGMT0_nodename'
```

Verify and Confirm the Software Upgrade

Use the following steps to verify and confirm that the upgrade script was successful:

The upgrade script will have stopped Lustre after completing the upgrade. Lustre will therefore be started again in this verification procedure.

22. Verify that the failback was successful:

```
[root@n000]# crm_mon -1
```

Sample output:

```
=====
Last updated: Thu Dec 25 08:45:30 2014
Last change: Thu Dec 25 06:54:56 2014 via crm_resource on snx11000n000
Stack: Heartbeat
Current DC: snx11000n001 (2d7734d6-19f3-4209-95d6-8b3091a33801) - partition
with quorum
Version: 1.1.6.1-6.el6-0c7312c689715e096b716419e2ebc12b57962052
2 Nodes configured, unknown expected votes
51 Resources configured.
=====
Online: [ snx11000n000 snx11000n001 ]
snx11000n000-1-ipmi-stonith (stonith:external/ipmi): Started snx11000n000
snx11000n001-1-ipmi-stonith (stonith:external/ipmi): Started snx11000n001
snx11000n000-2-ipmi-stonith (stonith:external/ipmi): Started snx11000n000
snx11000n001-2-ipmi-stonith (stonith:external/ipmi): Started snx11000n001
Clone Set: cln-diskmonitor [diskmonitor]
Started: [ snx11000n000 snx11000n001 ]
Clone Set: cln-last-stonith [last-stonith]
Started: [ snx11000n000 snx11000n001 ]
Clone Set: cln-kdump-stonith [kdump-stonith]
Started: [ snx11000n000 snx11000n001 ]
prm-httpd (lsb:httpd): Started snx11000n000
prm-mysqld (lsb:mysql): Started snx11000n000
prm-nfslock (lsb:nfslock): Started snx11000n001
prm-bebundd (lsb:bebundd): Started snx11000n000
Clone Set: cln-cerebrod [prm-cerebrod]
Started: [ snx11000n001 snx11000n000 ]
prm-conman (lsb:conman): Started snx11000n000
Clone Set: cln-syslogng [prm-syslogng]
Started: [ snx11000n001 snx11000n000 ]
prm-dnsmasq (lsb:dnsmasq): Started snx11000n001
prm-nodes-monitor (lsb:nodes-monitor): Started snx11000n000
Clone Set: cln-ses_mon [prm-ses_monitor]
Started: [ snx11000n001 snx11000n000 ]
Clone Set: cln-nsca_passive_checks [prm-nsca_passive_checks]
Started: [ snx11000n001 snx11000n000 ]
Resource Group: grp-icinga
prm-icinga (lsb:icinga): Started snx11000n000
prm-nsca (lsb:nsca): Started snx11000n000
prm-npcd (lsb:npcd): Started snx11000n000
prm-repo-local (ocf::heartbeat:Filesystem): Started snx11000n001
prm-repo-remote (ocf::heartbeat:Filesystem): Started snx11000n000
prm-db2puppet (ocf::heartbeat:oneshot): Started snx11000n000
Clone Set: cln-puppet [prm-puppet]
Started: [ snx11000n001 snx11000n000 ]
prm-nfsd (ocf::heartbeat:nfsserver): Started snx11000n001
prm-vip-eth0-mgmt (ocf::heartbeat:IPAddr2): Started snx11000n000
prm-vip-eth0-ipmi-sec (ocf::heartbeat:IPAddr2): Started snx11000n000
prm-vip-eth0-nfs (ocf::heartbeat:IPAddr2): Started snx11000n001
Resource Group: snx11000n000_md64-group
snx11000n000_md64-raid (ocf::heartbeat:XYRAID): Started snx11000n000
snx11000n000_md64-fsys (ocf::heartbeat:XYMNTR): Started snx11000n000
snx11000n000_md64-stop (ocf::heartbeat:XYSTOP): Started snx11000n000
Resource Group: snx11000n000_md67-group
```

```

snx11000n000_md67-raid (ocf::heartbeat:XYRAID): Started snx11000n001
snx11000n000_md67-fsys (ocf::heartbeat:XYMNTR): Started snx11000n001
snx11000n000_md67-stop (ocf::heartbeat:XYSTOP): Started snx11000n001
prn-ctdb (lsb:ctdb): Started snx11000n001
Resource Group: grp-plex
prn-rabbitmq (lsb:rabbitmq-server): Started snx11000n000
prn-plex (lsb:plex): Started snx11000n000
Resource Group: grp-dcs
prn-zabbix-server (lsb:zabbix-server): Started snx11000n000
prn-zabbix-listener (lsb:agent-listener): Started snx11000n000
baton (ocf::heartbeat:baton): Started snx11000n000
snx11000n000_mdadm_conf_regenerate (ocf::heartbeat:mdadm_conf_regenerate):
Started snx11000n000
snx11000n001_mdadm_conf_regenerate (ocf::heartbeat:mdadm_conf_regenerate):
Started snx11000n001

```

23. Confirm that all nodes are running the new software version:

```
[root@n000]# cscli show_node_versions -g all
```

Following is an example command and output showing all nodes at the 1.5.0 software version:

```

[root@snx11000n000 admin]# cscli show_node_versions -g all
snx11000n000 1.5.0-38
snx11000n001 1.5.0-38
snx11000n002 1.5.0-38
snx11000n003 1.5.0-38
snx11000n004 1.5.0-38
snx11000n005 1.5.0-38

```

24. Start Lustre using CCLI:

```
[root@n000]# cscli mount -f filesystem_name
```

Following is an example command and output:

```

[root@snx11000n000 admin]# cscli mount -f testfs
mount: MGS is starting...
mount: MGS is started!
mount: Database assertion: created a new connection but pool_size is already
reached (4 > 3)!
mount: starting testfs on snx11000n[002-403]...
mount: Database assertion: created a new connection but pool_size is already
reached (4 > 3)!
mount: starting testfs on snx11000n[006-407]...
mount: starting testfs on snx11000n[004-405]...
mount: No resources found on nodes snx11000n[000-401] for "testfs" file system
mount: starting testfs on snx11000n[008-409]...
mount: testfs is started on snx11000n[002-403]!
mount: testfs is started on snx11000n[008-409]!
mount: testfs is started on snx11000n[006-407]!
mount: testfs is started on snx11000n[004-405]!
mount: File system testfs is mounted.

```

25. Verify that Lustre started:

```
[root@n000]# cscli fs_info
```

Following is an example command and output:

```
[root@snx11000n000 admin]# cscli fs_info
```

```
-----
Information about "testfs" file system:
-----
```

| Node | Node type | Targets | Failover partner | Devices |
|--------------|-----------|---------|------------------|--|
| snx11000n002 | mgs | 1 / 1 | snx11000n003 | |
| snx11000n003 | mds | 1 / 1 | snx11000n002 | /dev/md66 |
| snx11000n008 | oss | 4 / 4 | snx11000n009 | /dev/md0, /dev/md2, /dev/ md4, /dev/md6 |
| snx11000n009 | oss | 4 / 4 | snx11000n008 | /dev/md1, /dev/md3, /dev/ md5, /dev/md7 |
| snx11000n006 | oss | 4 / 4 | snx11000n007 | /dev/md0, /dev/md2, /dev/ md4, /dev/md6 |
| snx11000n007 | oss | 4 / 4 | snx11000n006 | /dev/md1, /dev/md3, /dev/ md5, /dev/md7 |
| snx11000n004 | oss | 4 / 4 | snx11000n005 | /dev/md0, /dev/md2, /dev/ md4, /dev/md6 |
| snx11000n005 | oss | 4 / 4 | snx11000n004 | /dev/md1, /dev/md3, /dev/ md5, /dev/md7 |

26. Verify Lustre file system functionality by mounting a client and creating a file.

- a. Make a directory named /mnt/test and mount it:

```
[root@n000]# mkdir /mnt/test
[root@n000]# pdsh -g mgs=primary lctl list_nids
snx11000n002: 172.18.1.3@o2ib
[root@n000]# cscli fs_info | grep file
```

Information about "testfs" file system:

```
[root@n000]# mount -t lustre 172.18.1.3@o2ib:/testfs /mnt/test
[root@n000]# mount | grep /mnt/test
172.18.1.3@o2ib:/testfs on /mnt/test type lustre (rw)
[root@n000]# df -h
```

- b. In directory /mnt/test/, create a 1GB file, remove the file, and then remove the directory:

```
[root@n000]# cd /mnt/test/
[root@n000]# dd if=/dev/zero of=testdata bs=1M count=1024
[root@n000]# ls -la
[root@n000]# rm testdata
[root@n000]# cd /
[root@n000]# umount /mnt/test
[root@n000]# rmdir /mnt/test
```

Post-Installation Steps

Subtasks:

- [Apply 1.5.0 System Updates If Applicable](#) on page 28
- [Post-Installation Firmware Requirements](#) on page 28
- [Configure the Management Switches](#) on page 29
- [Configure Brocade Management Switches](#) on page 29
- [Configure NetGear Management Switches](#) on page 31

Apply 1.5.0 System Updates If Applicable

Once the Sonexion 1600/900 software has been upgraded to 1.5.0, apply any required system updates.

Perform this step even if the System Update was used to prepare the cluster before running the upgrade script. See [Prepare the Cluster with the Latest System Update for 1.5.0](#) for more information.

1. Identify any new system updates released since Sonexion 1600/900 1.5.0 became available.
2. Determine which system updates have not yet been applied to the system.
3. Download the system updates, identified in the previous step, from CrayPort.
4. Install the downloaded system updates per the system update release notes.

For the most up-to-date information about available system updates for Sonexion 1.5.0, check CrayPort. System Updates are cumulative, so look for the most recent SU number for 1.5.0.

Post-Installation Firmware Requirements

The following system firmware may need leveling after completing the software upgrade to Sonexion 1.5.0.

- Mellanox HCA firmware.
- LSI firmware
- USM firmware.
- Firmware running on the HDDs and SSDs
- Mellanox InfiniBand switch firmware..

Instructions to level system firmware can be found in separate standalone documents for each firmware. See [Reference Documents](#) on page 6 for a list of these documents.

The actual firmware package(s) needed may vary depending on installation requirements. If the firmware files are not available from CrayPort for the related firmware document, please contact Cray Support.

Configure the Management Switches

The Sonexion system's management switches (Brocade and NetGear) must be fully configured after the upgrade script has been completed. This configuration:

- Sets SSH parameters
- Sets the username and password
- Disables telnet access
- Save the proper Spanning Tree settings for all the ports on the system

The configuration is performed using scripts that are copied to the secondary MGMT node during the software upgrade process. The scripts are located in the `/mnt/nfsdata/tftpboot/sw_configs` directory.

Required Equipment:

Verify that the following equipment is available before configuring the management switches.

- Console with monitor and keyboard OR
PC/Laptop with a DB9 serial COM port, mouse, keyboard. The serial port should be configured for 9600 Kbps, 8 data bits, no parity and 1 stop bit.
- Rollover/Null modem serial cable (DB9 to RJ45 Ethernet)
- Serial terminal emulator program, such as SecureCRT, PuTTY, Tera Term, Terminal Window (Mac) etc.
- ESD strap, boots, garment or other approved methods

To configure the management switches:

- For Brocade switches, go to [Configure Brocade Management Switches](#) on page 29.
- For NetGear switches, go to [Configure NetGear Management Switches](#) on page 31.

Configure Brocade Management Switches

Perform the following procedure for each Brocade management switch in the system.

1. To connect to the switch, connect the rollover cable to the console port (RJ45) at the left side of the Management Switch and to the serial COM port (DB9) on the PC/Laptop. Refer to the following figure.

Figure 5. Console Port on Brocade ICX 6430 Switch



2. Open a terminal session. To use a PuTTY/SecureCRT emulator program, configure the serial port for 9600 Kbps, 8 data bits, no parity and 1 stop bit.

To use Terminal Window on a Macintosh, type the following and then press **Enter**:

```
screen /dev/ttyUSB0 cs8,9600
```

After connecting to the switch, the terminal screen should display the switch command prompt, which is shown as `[switch]` in the remaining commands.

3. Verify that there is an IP address assigned to the switch:

```
[switch]show ip address
```

If the command output shows that there is no IP address assigned to the switch, perform these steps, one at a time, correct any identified problems, and then check for an IP address again. When the switch has an IP address, go to the following step.

- a. Verify the switch's connections and cabling.
- b. Verify that the Sonexion system is powered up.
- c. Verify the switch's DHCP configuration. Run these commands:

```
[switch]enable
[switch]config t
[switch]ip dhcp-client enable
[switch]ip dhcp-client auto-update enable
[switch]write mem
```

4. Load the appropriate configuration file from the secondary MGMT node to the switch.
 - a. Determine which configuration file is needed from the following table.

Table 6. Brocade Configuration File Selection

| Rack | Switch | Switch Type | Configuration File to Use |
|---------|-------------------------|-------------|---------------------------|
| Base | Primary (Lower - SW0) | 24 ports | Brocade24_Pri.cfg |
| Base | Secondary (Upper - SW1) | 24 ports | Brocade24_Sec.cfg |
| Base | Primary (Lower - SW0) | 48 ports | Brocade48_Pri.cfg |
| Base | Secondary (Upper - SW1) | 48 ports | Brocade48_Sec.cfg |
| Storage | Primary or Secondary | 24 ports | Brocade24_Stor.cfg |
| Storage | Primary or Secondary | 48 ports | Brocade48_Stor.cfg |

- b. Load the configuration file to the switch. Run the following commands, where `<config_file_name>` is the name of the appropriate configuration file from the preceding table:

```
[switch]enable
[switch]copy tftp start 172.16.2.4 /sw_configs/config_file_name
[switch]reload
```

5. Disconnect the PC or terminal from the switch.
6. Repeat Steps 1 through 5 for each Brocade management switch in the system.

Configure NetGear Management Switches

Perform the following procedure for each NetGear management switch in the system.

1. To connect to the switch, connect the rollover cable to the console port (RJ45) at the left side of the Management Switch and to the serial COM port (DB9) on the PC/Laptop.
2. To use a PuTTY/SecureCRT emulator program, open a terminal session and specify these serial port settings:

| | |
|-------------------------------------|-------------|
| Baud rate (bits per second): | 9600 |
| Data bits: | 8 |
| Stop bits: | 1 |
| Parity: | None |
| Flow control: | Off |

To use Terminal Window on a Macintosh, type the following and press **Enter**:

```
screen /dev/ttyUSB0 cs8,9600
```

After connecting to the switch, the terminal screen should display the switch command prompt, which is shown as `[switch]` in the remaining commands.

3. Load the configuration file from the secondary MGMT node to the switch. Run these commands on the switch:

```
[switch]#enable
[switch]#copy scp://admin@172.16.2.4/mnt/nfsdata/tftpboot/sw_configs/
netgear.cfg nvram:startup-config
```

When prompted for "Remote Password", enter the Sonexion admin password of the management node. When prompted "Are you sure you want to start? (y/n)", enter **y**.

```
[switch]#reload
```

Following is an example of these commands and the expected output:

```
(M5300-28G3)#enable
(M5300-28G3)#copy scp://admin@172.16.2.4/mnt/nfsdata/tftpboot/switch_configs/
netgear.cfg nvram:startup-config
Remote Password:*****
Mode..... SCP
Set Server IP..... 172.16.2.4
Path..... /mnt/nfsdata/tftpboot/
switch_configs
Filename..... netgear.cfg
Data Type..... Text Configuration
Download configuration file. Configuration will be applied upon next reboot.
Management access will be blocked for the duration of the transfer
Are you sure you want to start? (y/n) y
File transfer successful...
(M5300-28G3)#reload
```

4. Disconnect the PC or terminal from the switch.

5. Repeat Steps 1 through 4 on all NetGear management switches in the system.
6. The procedure to upgrade the Sonexion 1600/900 software to 1.5.0 is complete.

Troubleshoot from Error Messages

Errors depend on the environment; not the script. For example, if a node shuts down, the impact is on the HA, but the script did not provide the instruction for the node to shut down.

Additional error information may be collected from the Test Group, including workarounds. The list below can be expanded as needed.

This section describes possible problems seen with Sonexion 1600 or Sonexion 900 while the 1.5.0 upgrade script is running. If an error occurs, check the following list of known errors. The boldface text is the actual full error message. Description and Workaround comments have been added.

"Early error" means no changes have been introduced before corrections have been made by the script. If you get an "early error" message, fix the problem and rerun the script.

Can't mapfile \$yaml

Description: Can't mapfile \$yaml. Early error -- means Sonexion is really broken.

Workaround: STOP. Contact Cray Technical Support.

Can't patch beUSBUpdater

Description: Can't patch beUSBUpdater, see above for details. This usually means that the script was interrupted using "quit" in some late stages. Recovery depends on the phase it happens in. It's hard to tell now how to recover. It's better to contact support in that case.

Workaround: Each or any combination of the errors listed before this one may contribute to the "beUSBUpdater" problem. There are too many variables to describe a reasonable workaround.

STOP. Contact Cray Technical Support.

Root user

Description: You have to be a root user to run \$0. Early error -- Recovery: run "sudo su -"

Workaround: Change to root user:

```
[admin@n000]$ sudo su -[root@n000]#
```

Run script on primary management node

Description: Please run \$0 on \$mgmt_primary" - the script being executed is not on the primary management node. This indicates that the MGMT resources are not balanced.

Workaround: Rebalance the HA resources as described in Section 2.8.5 Run Failback before Starting the Upgrade Script, Step 2.

MGMT resources not on primary, need failback node

Description: MGMT resources are not running on \$mgmt_primary, please failback that node. This error also indicates that the MGMT resources are not balanced.

Workaround: Rebalance the HA resources as described in Section 2.8.5 Run Failback before Starting the Upgrade Script, Step 2.

Offline nodes are not available

Description: Nodes \$offline_nodes are unavailable.

Workaround: Please make sure all nodes accept login via SSH. Early error -- this is an improved message compared to the previous "not all are available"

To verify nodes are accepting SSH commands, run the command and look for a date prompt returned from each node "pdsh -date"

```
snx11000n000: Tue Sep 10 07:31:25 PDT 2013
snx11000n001: Tue Sep 10 07:31:25 PDT 2013
snx11000n002: Tue Sep 10 07:31:25 PDT 2013
snx11000n005: Tue Sep 10 07:31:25 PDT 2013
snx11000n004: Tue Sep 10 07:31:25 PDT 2013
snx11000n003: Tue Sep 10 07:31:25 PDT 2013
[root@snx11000n0-oem ~]#
```