

- [Cluster support](#)
- [High Performance Switch](#)

IBM HPS POWER5 Readme - Service Pack 17

Introduction

Anchor link navigation

- **[Introduction](#)**
- [Supporting documentation](#)
- [Component updates/download information](#)
- [Upgrade / Installation sequences](#)

- [AIX LP](#)
- [Known problems and workarounds](#)
- [Readme archive](#)

Overview

06/18/2008

The IBM pSeries High Performance Switch (HPS) POWER5 Service Pack provides updates and code levels for these components:

- Hardware Management Console (HMC).
- Global Firmware (GFW).
- Power Subsystem Microcode (BPC code).
- HPS Network Manager (HPSNM).
- HPS/Switch Network Interface (SNI) Devices LP.
- Cluster Systems Management (CSM) LP.
- Reliable Scalable Cluster Technology (RSCT) software.
- AIX 5L Version 5.2 and 5.3.
- Various supporting AIX LPs.

HPS POWER5 Service Pack 17 provides support for service updates for AIX5.2 TL10 (SP5), the new release of AIX5.3 (TL8 SP1), new mod levels and service updates for cluster / HPC components RSCT, CSM, PE, LAPI, LoadL and service updates for cluster / HPC components GPFS, VSD and the HPS

drivers.

Note, there are no new POWER5 firmware, HMC, or HPSNM updates between Service Pack 16 and Service Pack 17.

If you are coming from HPS Service Pack 15 or earlier, consider the following:

Issues moving to BASE GA7 that are resolved in Release Level SF240_219 are that 9118-575 systems with 01SF230_126_120 (or a later 01SF230 firmware level) that are upgraded to 01SF235_160_160 firmware will not boot if the I/O chassis on 9118-575 has been replaced. The failure requires that the VPD card for the 9118-575 be replaced to fix the problem.

If you are already at System Firmware GA7 SP2 (01SF240_233_201) or later, there is no new microcode. This means that if you are following the new code load procedure; then you will not have to re-IPL the CEC nor Power Off the CEC; you will only have to reboot the LPARs to activate the alt_disk.

This document also contains general guidelines for upgrading the components listed in [Component updates/Download information](#).

These guidelines are intended to be a supplement to the other IBM documents referred to in this document.

You must have all referenced documents available before you begin the upgrade process. A list of referenced documents can be found in [Supporting Documentation](#).

The Code Levels listed in [Component updates/Download information](#) reflect the levels available at the time of this HPS POWER5 Release. Some components support only a single version, notably the Microcode for GFW and the Power Subsystem. Subsequent released versions are expected to be backward compatible.

The procedure outlined in [Recommended Installation Sequence](#) is the standard sequence of installation. Non-standard sequences or undocumented code levels may cause unforeseen problems. In this event please contact your Customer Service Representative.

HPS POWER5 SP3 introduced a modified install plan intended to reduce the maintenance window down time. Refer to the [Alternate procedure to reduce downtime](#).

HPS POWER5 SP5 introduced support for system firmware level 240.

HPS POWER5 SP11 allows use of the HPS TOD "Time of Day" function (refer to "Restrictions" section).

Service Pack Levels

230 Service pack code levels:

HMC: Version: 4; Release: 5.0 - Build level 20050629 + MH00454
IBMhsc.HPSNM_log-1.4.1.3-1 RPM

CEC: Global Firmware (GFW) - Version SF230_158
FRAME:Power Code (BPC code) - Version BP230_155

235 Service Pack code levels:

HMC: Version 5.1 + MH00593
IBMhsc.HPSNM_log-1.4.1.1-1 RPM
IBMhsc.NM_Common-1.0.0.0-12 **CEC:** Global Firmware (GFW) - Version SF235_214

FRAME: Power Code (BPC code) - Version BP235_212

240 Service pack code levels:

HMC: Version: 6 Release: 1.0 + MH01012 (HMC V6 R1.3) + efix MH01082 (replaces MH01064)

HPSNM Service Pack 10: IBMhsc.HPSNM_log-1.4.1.12-1 RPM
IBMhsc.NM_Common-1.1.0.4-1 RPM
IBMhsc.IBMNM_GUI-1.1.0.3-1 RPM

CEC: Global Firmware (GFW) - Version SF 240_338
FRAME: Power Code (BPC code) - Version BP240_219

AIX Service Packs and Supported Cluster and High Performance Computing (HPC) LPs:

- AIX Version 5.2 TL 10 SP5.
- AIX Version 5.3 TL 08 SP1.
- HPSNM - csm.hpsnm 1.4.1.18 (CSM Management Server only).
- AIX 5L Version 5.2 / HPS Devices LP 1.1.3.x as listed in [Detailed LP level check](#).
- AIX 5L Version 5.3 / HPS Devices LP 1.2.0.x as listed in [Detailed LP level check](#).
- Other supported AIX LP's as listed in [Detailed LP level check](#).

New functions and features

The following items are new functions and features in release 240 for the High End systems that support

the IBM HPS:

- Support for 2.2 GHz processors on the model 570.
- Support for 1.9 GHz processor cards on the model 9406-570.
- Support for 4, 8, and 16 GB memory cards with 0% initial activation, and 1 GB activation increments, on model 590 and model 595 systems.
- Support for two CUoD DDR2 memory features on the model 570: a 4/8 GB feature, and an 8/16 GB memory feature.
- Support for mixing CoD-capable DDR2 memory with DDR2 memory that is not CoD-capable in a system.
- Support for F/C 0649, a new high-performance SCSI with RAID 6 disk controller.
- Support for the collection (and viewing on the HMC) of logical partitions' utilization of processor and memory resources.
- Support for enhanced model 575.
- Support for concurrent maintenance of the following entities on model 590 and 595 systems:
 - Node addition or upgrade.
 - RIO bus adapter addition or upgrade.
 - RIO bus adapter repair.
 - Clock card repair.
 - Various enhancements to the server firmware to reduce IPL times.
 - Support for huge pages (16 GB) in the Advanced System Management Interface (ASMI) menus.
 - Enhancements to the "Restore to factory default" option, CoD options, time-of-day, and firmware update policy menu on the ASMI menus.
 - Enhancements to the memory deconfiguration menu interface in the ASMI menus.
 - The option to set the number of virtual LAN (VLAN) switches was added to the ASMI menus.
 - A feature was added to the system firmware so that booting a system or partition from a hard disk will be retried, instead of immediately failing, to handle the situation in which disks are not yet available to boot before the firmware wants to boot from them. This typically happens when the boot disk is in an external disk enclosure.
 - Add support for USB devices in system firmware.
 - The maximum number of I/O towers on a loop was increased to six on iSeries model 595.

Health check/planning help

- [FLRT: Fix Level Recommendation Tool](#)
- [IBM Systems compatibility information](#)

- [Cluster support](#)
- [High Performance Switch](#)

IBM HPS POWER5 Readme - Service Pack 17

Supporting documentation

Anchor link navigation

- [Introduction](#)
- **[Supporting documentation](#)**
- [Component updates/download information](#)
- [Upgrade / Installation sequences](#)

- [AIX LP](#)
- [Known problems and workarounds](#)
- [Readme archive](#)

Overview

Use the following information to reference the supporting documentation for this Readme:

- [High Performance Switch Planning, Installation, and Service for IBM server p5 servers](#)
- [Hardware Management Console](#)
- [eServer pSeries and AIX Information Center](#)
- [Switch Network Interface for eServer pSeries High Performance Switch Guide and Reference](#)
- [AIX 5.2 documentation](#)
- [AIX 5L Version 5.2 Installation Guide and Reference](#)
- [AIX 5L Version 5.2 Performance Management Guide](#)
- [Reliable Scalable Cluster Technology \(RSCT\) Library](#)
- [General Parallel File System \(GPFS\) Library](#)
- [GPFS FAQs](#)
- [Cluster System Management \(CSM\) Library](#)
- [IBM Parallel Environment \(PE\) Library](#)
- [LoadLeveler for AIX5L and Linux V3.3](#)
- [Guide to Multiple Page Size Support on AIX 5L Version 5.3](#)
- [pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide for IBM](#)

[eServer p5 servers](#)

- [CSM for AIX5L V1.x Planning and Installation Guide](#)
- [CSM for AIX5L V1.x Administration Guide](#)

Health check/planning help

- [FLRT: Fix Level Recommendation Tool](#)
- [IBM Systems compatibility information](#)

- [Cluster support](#)
- [High Performance Switch](#)

IBM HPS POWER5 Readme - Service Pack 17

Component updates / Download information

Anchor link navigation

- [Introduction](#)
- [Supporting documentation](#)
- **[Component updates/download information](#)**
- [Upgrade / Installation sequences](#)
- [AIX LP](#)
- [Known problems and workarounds](#)
- [Readme archive](#)

Component Updates/Download Information

The following table contains information about component updates and locations for downloading the updates.

Component	Service pack updates	Download sites
HMC This release requires that the HMC is installed with code levels that are related to specific firmware.	230 HMC V4R5.0. Select HMC Version 4.5 HMC_Update_V4R5.0_1.zip HMC_Update_V4R5.0_2.zip This release requires efix MH00454. You can download the HMC Recovery CD V4R5.0 package as a set of .iso images by using Download Director or you can download the corrective service file directly from the Internet.	HMC Support Select "HMC 4.5" Downloads for updates and fixes are on the "Corrective service" tab. For the HMC Update ReadMe, select "View" in the "ReadMe" column for the specific update package.

See Notice regarding 7310-C04 and 7310-C05/
internal modem compatibility issue.

For fix locations and
installation instructions,
select the "Installation
instructions" link on the
"Corrective service" tab.

235

HMC V5R1

Select "HMC_Recovery_CD_V5R1.0"

[HMC Support](#)

This release requires that the HMC is installed with
HMC V5R1.0.

Select "HMC 5.1"

PTF MH00593 is an HMC Required Maintenance
Package (HMC Version 5, Release 1.2.). This
Maintenance Fix *must* be installed on HMC Version
5 Release 1.0.

Downloads and media
ordering for Recovery
DVDs are on the
"Recovery media" tab.
Click the "View" link for
the ISO images to see the
HMC Readme.

Updating to HMC 5.1 requires an Upgrade via the
Recovery media process. There is no "update"
corrective service to upgrade your current HMC
Version to HMC 5.1.

Downloads for updates
and fixes are on the
"Corrective service" tab.

The HMC Recovery DVD V5 R1.0 is a bootable
image and contains the HMC Base Code. You can
order the package DVDs from the HMC "Recovery
Media" tab. You can also download the DVD images
in ISO format, and then use the images to burn your
own DVD.

For the HMC Update
ReadMe, select "View" in
the "ReadMe" column for
the specific update
package.

Click the "Installation instructions" tab for
procedures for burning DVDs and for using the
HMC Recovery DVD V5 R1.0.

For fix locations and
installation instructions,
select the "Installation
instructions" link on the
"Corrective service" tab.

See Notice regarding 7310-C04 and 7310-C05/
internal modem compatibility issue.

240

HMC V6R1.0 + MH01012 (updates HMC V6R1.0, V6R1.1 or V6R1.2 to V6R1.3)

This release requires that the HMC is installed with HMC V6R1.3.

See Notice regarding 7310-C04 and 7310-C05/ internal modem compatibility issue.

Notes:

1. MH01012 is required to get to V6R1.3.
2. You must install HMC Version 6 Release 1.3 before installing efix MH01082 (refer to the HMC V6.1.x readme for information about this efix.)
3. You must use Recovery media or the Network to perform an Upgrade to update to HMC Version 6 Release 1.0.

[HMC Support](#)

Under "HMC products for servers with POWER5 processors" and select "HMC 6.1.3"

Downloads for updates and fixes are on the "Corrective service" tab.

For the HMC Update ReadMe, select "View" in the "ReadMe" column for the specific update package.

For fix locations and installation instructions, select the "Installation instructions" link on the "Corrective service" tab.

Contact your IBM Sales Representative or Business Partner, and order Hardware Feature Code (MES) 0960 for the initial upgrade CDs.

230: Fix pack BP230_155
Fix pack SF230_158

235: Fix pack BP235_212**
Fix pack SF235_214

**Note:

Power Subsystem update BP235_212 should only be installed on POWER5 systems using the IBM POWER Systems High Performance Switch.

Download this update by clicking on the following links:

[Microcode Updates](#)

AND

- [02bp235_212_156.rpm](#)

System Firmware

- [02bp235_212_156.xml](#)
- [02bp235_212_156.html](#)

240: Fix pack BP240_219

Fix pack SF240_338

Power Subsystem Microcode is available up to 1 week earlier and GFW is available 3-4 days earlier than the web download site.

Please contact your IBM Service Support Rep (SSR) for the Power Subsystem Microcode or GFW package if not available on website.

AIX levels for this Service Pack:

AIX52

Service Pack 5200-10-05-0819

bos.mp64: 5.2.0.112

[AIX V 5.2](#)[AIX V 5.3](#)**AIX****AIX53**

Service Pack 5300-08-01-819

bos.mp64: 5.3.8.1

Note:

Because critical and pervasive fixes are now delivered as part of Service Packs, Critical Fix Packs will no longer be created. To better understand this direction, read [AIX Service Strategy and Best Practices](#).

Notes:

1. CSM 1.4 requires RPM update openCIMOM 0.8(5.2)
2. If you are APPLYING software for csm. server 1.4.1.1, please install the following images from the AIX Installation Media Volume 2:

tcl**tk****expect**

openCIMOM
(CSM requisite)

[openCIMOM](#)

conserver-8.1

SSL_
(CSM requisite)

openssl-0.9.7d-2

sourceforge.net - openssh-aix

AIX-rpm-5.2.0.40-1

SSH_
(CSM requisite)

openssh-3.8.1p1_52.tar.gz

sourceforge.net - openssh-aix

On the CSM MS:

This release requires AIX HPSNM LP csm.hpsnm 1.4.1.18

apply HPSNM AIX PTF U812753, APAR IY96115

[Fix Central](#)

Note:

AIX filesets csm.hpsnm and csm.server may be included in AIX Maintenance or Technical Level Update Packages, should *only* be installed on the CSM MS and *not* on production LPARS.

On the HMC:

HPSNM

230 Level:

HMC v 4.5

HPSNM Service Pack 3 MB01192 MH00401

HPSNM_Serv_Pack3.zip

235 Level:

HMC v 5.1

HPSNM Service Pack 4 MB01237 MH00461

HPSNM_Serv_Pack4.zip

240 Level:

HMC v 6.1.2

HPSNM Service Pack 10 MB02132 MH01013

HPSNM_Serv_Pack10.zip

[Hardware Management Console](#)

Support for UNIX servers and Midrange servers
Select an HMC release, and then select the "HPSNM/IBMNM fixes" tab

Health check/planning help

- [FLRT: Fix Level Recommendation Tool](#)
- [IBM Systems compatibility information](#)

- [Cluster support](#)
- [High Performance Switch](#)

IBM HPS POWER5 Readme - Service Pack 17

Upgrade/Installation sequences

Anchor link navigation

- [Introduction](#)
- [Supporting documentation](#)
- [Component updates/download information](#)
- **[Upgrade / Installation sequences](#)**

- [AIX LP](#)
- [Known problems and workarounds](#)
- [Readme archive](#)

Upgrade procedure

Perform the following upgrade procedure to reduce the maintenance window for the service packs upgrades.

1. On the HMC - install / upgrade:
 - Upgrade ELA master and wait until ELA master is up and running.
 - Install Primary update.
 - Install miscellaneous updates.
 - Install HPSNM updates.
 - Upgrade the rest of the HMCs.
2. On the CSM MS - install / upgrade:
 - Install the new AIX PTFs.
 - Install the new RSCT PTFs.
 - Install the new CSM PTFs.
 - Install the new FNM PTFs.
 - Reboot the CSM MS.
3. Power Code upgrade:
 - If the new Power Code supports concurrent upgrade; then select concurrent upgrade

- during the update process.
 - If the new Power Code does not support concurrent upgrade; then select concurrent install only with deferred disruptive activate.
 - Update all node frames and switch frames.
4. GFW upgrade:
 - Disable SNM (only on disruptive upgrades).
 - If the new GFW supports concurrent upgrade; then select concurrent upgrade during the update process.
 - If the new GFW does not support concurrent upgrade; then select concurrent install only with deferred disruptive activate.
 - Enable SNM.
 5. On the LPARs - install / upgrade AIX and LPPs:
 - Use alternative disk install for AIX and LPPs.
 - Change the bootlist to the alternative disk.
 6. Stop the workload:
 - If you do not have to activate Power Code or GFW; then reboot the LPARs to activate alt disk install.
 - If you do not have to activate Power Code or GFW but have new microcode; then shutdown the LPARs.
 - Re-IPL CECs.
 - Activate the LPARs.
 - If you do have to activate Power Code or GFW:
 - Shutdown the LPARs.
 - Power off the CECs.
 - Activate the Power Code or the GFW Code.
 - Power on The CEC(s) the LPARs.
 7. Restart the workload.

This completes the Upgrade Sequence.

Hint: To decrease installation time for an upgrade, you may choose to use a globally mounted General Parallel File System (GPFS) over a high speed switch network (ie. HPS Switch). The following is an alternative to using the alternate-disk install method in steps 5 and 6 above:

1. Do alternate-disk install or

Copy the AIX and LP file set upgrade images to a local file system on each LPAR. You can do this by putting the images first in a GPFS and then copy them to each LPAR.

2. Stop the workload.

Upgrade the AIX and Cluster LPs on the LPARs using the dsh command (ie: /scratch is a local file system on each LPAR and it contains the AIX and LPP file sets) From the Management

Server, run these commands to upgrade the LPARs:

```
dsh -av "lspp -l -J > /tmp/file.aix"
```

```
dsh -av "echo installp -e /tmp/inst.log -aXYgd /scratch -f /tmp/file.aix | at now"
```

Continue with the rest of Step 6 about rebooting the LPAR or power cycle to activate, as appropriate.

[Back to top](#)

Install procedure

Complete the information in the following subsections to install this service pack:

1. [Install HMC.](#)
2. [Install Power Subsystem Microcode.](#)
3. [Install GFW.](#)
4. [Create LPAR definition and assign adapters.](#)
5. [Install Cluster System Manager \(CSM\) Software.](#)
6. [Define the nodes/LPARs in the cluster.](#)
7. [Assign IP addresses for SNIs.](#)
8. [Install AIX on all LPARs in the cluster.](#)
9. [Apply HPS/SNI updates to all LPARs.](#)
10. [Dynamic Host Configuration Protocol \(DHCP\).](#)
11. [Power Off CECs.](#)
12. [Configure the High Performance Switch Network Manager \(HPSNM\).](#)

In order to coordinate the tasks necessary for the installation and configuration of the HPS Cluster, we recommend that Service Support Representatives and System Administrators review the following documents applicable to your CSM release:

- [pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide for IBM eServer p5 servers](#)
- [CSM for AIX5L V1.x Planning and Installation Guide](#)
- [CSM for AIX5L V1.x Administration Guide](#)

Especially Chapter 6 *Installation* of the [pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide for IBM eServer p5 servers](#) provides detailed information on the hardware, firmware, and software sequence of installation.

Step 1: Install HMC

Complete the following steps to install the Hardware Management Console (HMC) software on all HMCs in the cluster:

1. [HMC Software Download Information](#)
2. [Installation Notes](#)
3. [Level Check](#)
4. [Level Check from HMC GUI](#)
5. [Service Network](#)
6. [WEBSM](#)
7. [HMC Properties](#)

There are no required fixes to apply.

Use the following documentation for reference:

- [Hardware Management Console](#)
- [Hardware Management Console V4R5.0 Readme](#)
- [Hardware Management Console V5R1.0 Readme](#)
- [Hardware Management Console V6R1.x Readme](#)
- [Hardware Management Console for pSeries Maintenance Guide](#)
- [Hardware Management Console for pSeries Installation and Operations Guide](#)
- [pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide for IBM eServer p5 servers \(SA38-0646-00\)](#)

HMC Software Download Information

In the HMC corrective service support section on the [HMC Power5 servers page](#), scroll to "Supported HMC products" and select from the following:

- SF230: Version 4.5 machine code updates
- SF235: Version 5.1 machine code updates
- SF240: Version 6.1 machine code updates

Follow the information and HMC installation instructions.

Note: Check if your HMC is listed in the BIOS updates sub-section. If you are updating the HMC on a listed model, you must first update the BIOS of that HMC model. The BIOS and installation instructions are also available on the [HMC support page](#).

Installation Notes

If the HMC was used in standalone mode before CRHS setup, reboot the HMC after running the `addpeer` command to add the HMC into the peer domain. This ensures that the HMC GUI displays hardware in cluster mode.

Upgrading an HMC in the ClusterPeerDomain should not require that the `addpeer` command be run again for that HMC. The ClusterPeerDomain should be preserved after an upgrade.

Caution

There is a known problem when upgrading the HMC from Version 5 and early Version 6 to Version 6.1 or later. In this scenario the hardware server resource manager on the HMC fails to start after the upgrade. This results in the HMC not displaying the hardware configuration.

To avoid this problem when upgrading to HMC version 6.1.2 (in this p5 HPS SP13 Service Pack release), you must also install HMC interim fix PTF MH00946. Obtain this fix from the HMC 6.1.2 [HMC support page](#). MH00946 will be picked up in a future HMC V6.x update package.

To preserve the Cluster Ready Hardware Server configuration on an HMC, be sure to save "Save upgrade data" prior to upgrading the HMC. This data should automatically be restored after the HMC is upgraded and doesn't require a reboot of the HMC.

Install the HMC code using the most recent copy of the [HPS Install and Planning Guide](#).

Review the HMC information on the web page where you downloaded the images.

For new HMC install follow the instructions in Chapter 3, "Installing and Configuring the HMC" ["Hardware Management Console for pSeries Installation and Operations Guide"](#).

Level Check

Level Check by running the `lshmc -V` command on the HMC command line or selecting the "Help --> About Hardware Management Console" option on the HMC GUI.

Level Check from HMC command line where `lshmc -V` shows:

SF230:

version= Version: 4

Release: 5.0

HMC Build level 20051010.1

MH00454: Maintenance Package for V4R5.0 (10-10-2005)

SF235:

version= Version: 5

Release: 1.1

HMC Build level 20060104.1

SF240:

version= Version: 6

Release: 1.3

HMC Build level 20070910.1

Level Check from HMC GUI

Menu "Help" "About Hardware Management Console" shows:

230 level:

Version: 4

Release: 5.0

HMC Build level 20051010.1

MH00454: Required Maintenance Fix for V4R5.0 (10-10-2005)

235 level:

Version= Version: 5

Release: 1.0

HMC Build level 20060104.1

MH00523: Remote HMC Management fix (01-12-2006)

MH00607: Fix Dump collection issues (02-25-2006)

240 level:

Version= Version: 6

Release: 1.3

HMC Build level 20070910.1

Service Network

The HMC requires a 'Service Network' to communicate with the Cluster components:

- HMC
- Managed Frames
- Managed Systems (CEC FSP's)
- CSM Management Server (MS)

This service network is an Ethernet hub to which the cluster components are connected via an Ethernet cable. There is a requirement to connect only one port of the frames A-side BPA Ethernet hub to the primary service network (a.k.a. external Ethernet hub).

Connect an Ethernet cable from the A-side BPA network hub at the top of the frame (labeled J00 RJ45 ports A, B, C or D). All 24-inch frames, including frames for the p575, p590, p595 servers and switch-only frame, are provided with their internal networks pre-cabled to the internal network hub on both sides of the frame. No additional internal Ethernet cable connections are required.

The other ports on the hub are connected to frame components such as the SP and BPA. HMCs typically have two integrated Ethernet ports labeled 1 and 2, eth0 and eth1, respectively. eth0 must be connected to the private service network. Any other Ethernet port can be used for a public network if desired.

Check if BIOS hyperthreading is disabled on the HMC. Many of the rack mounted HMCs (8187-KUH, 7315-C03) have a BIOS option to enable hyperthreads. The HMC's embedded kernel will not run well when this option is enabled. Hyperthreading must be disabled before upgrading to HMC R4V5 or higher.

WEBSM

After a successful HMC installation or upgrade, the Install Shield version of the Web-based System Manager Remote Client a.k.a WebSM PC Client may need to be re-installed on your remote server or PC. Uninstall any existing WebSM PC Client.

Failure to uninstall will result in undefined behavior of the WebSM PC Client.

To install the PC Client software on your remote server or PC, complete the following steps:

1. Go to http://<hmc-hostname>/remote_client.html.
2. Select the image to install: Linux Or Windows.
3. Select "Open" to immediately install on the machine on which the browser is running.

Note:

The installation may take up to 10 minutes to complete.

For complete instructions on installing and using the Remote Client please refer to Chapter 9, "Installing and Using the Remote Client" in the [Hardware Management Console for pSeries Installation and Operations Guide](#).

HMC Properties

Additional information on the HMC properties is also available in Chapter 6, "Installation" of the [pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide for IBM eServer p5 servers](#).

Important: The login available at virtual console 0 (via the CTRL-ALT-F1 key sequence) is no longer available.

Set the date and time

Check that the date and time is set on the Hardware Management Console. Refer to "Step 3: Set the date and time on the console" of the "Firmware Installation" section in Chapter 6 ("Installation") of the [pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide for IBM eServer p5 servers \(SA38-0646-00\)](#) for more details.

Enable the CRHS port

Enable the Cluster Ready Hardware Server 8899:tcp port on the service network on the HMCs in the cluster.

To enable the Cluster Ready Hardware Server port, complete the following steps:

1. Open the HMC GUI.
2. Select *Customize Network Configuration*.
3. Select *Ethernet interface*.
4. Click on *Details*.
5. Click on *Firewall*.
6. Highlight *Cluster Ready Hardware Server 8899:tcp* from upper window.
7. Click the *Allow incoming* button.
8. Verify that *Cluster Ready Hardware Server* appears in the lower part of the window.
9. Reboot HMC after the configuration.

Set frame numbers

Set frame numbers on each frame in the cluster. Refer to Step 4: Setting frame numbers" of the "Firmware Installation" section in Chapter 6 ("Installation") of the [pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide for IBM eServer p5 servers](#) for more details.

[Back to top](#)

Step 2: Install Power Subsystem Microcode

Install Power Subsystem Microcode on all server and switch frames in the cluster.

Refer to "Step 5: Installing power subsystem microcode and managed system (GFW) firmware" of the "Firmware Installation" section in Chapter 6 ("Installation") of the pSeries High Performance Switch (HPS) Planning, Installation and for IBM eServer p5 server (SA38-0646-00) for installation details.

Important Preliminary Notes:

- A complete install/update consists of two files, (i) a firmware code fixpack in RPM format and (ii) a cover letter in XML format.
- The managed frames must be in the Standby/Standby state on the "Frame Management" panel prior to the power subsystem microcode install/upgrade.

Documentation Reference:

- [pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide for IBM eServer p5 servers](#)

Power subsystem microcode Download Information:

- <http://www14.software.ibm.com/webapp/set2/firmware/gjsn>

Level Check

To view the expected installed and activated power subsystem licensed internal code (LIC) level after a successful install/upgrade:

- On the HMC GUI, select the "Licensed Internal Code Updates" tab.
- From the "Change Internal Code" task panel, select "View system information" and click OK.
- Select "None" on the "Specify LIC Repository" panel and click OK.
- EC Number LIC Type Machine Type/Model/Serial Number Installed Level Activated Level Accepted Level: where <nnn> is the Accepted LIC level last successfully installed and activated power subsystem LIC level.

Refer to the [Component Update/Download Information](#) section for the Power Subsystem Microcode.

Note:

The HMC command line can also be used to update power subsystem microcode using the `lslic` and `updlic` commands. Refer to the HMC man pages for detailed information on using using these commands.

[Back to top](#)

Step 3: Install GFW

Install managed system firmware on all managed systems/CECs in the cluster. Refer to "Installing power subsystem microcode and managed system (GFW) firmware" in the "Firmware Installation" section in Chapter 6 ("Installation") of the [pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide for IBM eServer p5 servers](#) for installation details.

Important Preliminary Notes:

- A complete install/update consists of two files, (i) a firmware code fixpack in RPM format and (ii) a cover letter in XML format.
- The managed systems must be in either Operating, Standby, or Power Off state on the "Server Management" GUI panel prior to the managed system firmware install/upgrade.
- Verify that the managed frames are in the Standby/Standby state on the "Frame Management" GUI panel prior to the power subsystem microcode install/upgrade.

Documentation Reference:

- [pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide for IBM eServer p5 servers](#)

Managed system firmware (GFW) Download Information:

- <http://www14.software.ibm.com/webapp/set2/firmware/gjsn>

Level Check

To view the expected installed and activated power subsystem licensed internal code (LIC) level after a successful install/upgrade:

- On the HMC GUI, select the *Licensed Internal Code Updates* tab.
- On the *Change Internal Code* task panel, select *View system information* and click *OK*.
- Select "None" on the "Specify LIC Repository" panel and click OK.
- EC Number LIC Type Machine Type/Model/Serial Number Installed Level Activated Level Accepted Level: [See: GFW Levels](#)

where <nnn> is the Accepted LIC level last successfully installed and activated power subsystem LIC level

Note:

The HMC command line can also be used to update power subsystem microcode using the `lslic` and `updlic` commands. Refer to the HMC man pages for detailed information on using using these commands.

[Back to top](#)

Step 4: Create LPAR definition and assign adapters

For details refer to "HMC code load" in the "Software Installation" section in Chapter 6, "Installation", [pSeries HPS Planning, Installation and Service Guide](#).

[Back to top](#)

Step 5: Install Cluster System Manager Software

CSM software is part of the base AIX installation. The recommended AIX service levels: [See Component Updates](#).

a. Preliminary Notes:

- CSM requires openCIMOM 0.8(5.2) RPM update which is available on the AIX Toolbox page: <http://www.ibm.com/servers/aix/products/aixos/linux/download.html>
- The following opensource images also available on the AIX Toolbox page, <http://www.ibm.com/servers/aix/products/aixos/linux/download.html>
- pre-requisites for applying csm.server which will be carried out in Step 6b below:
 - tcl
 - tk
 - expect
 - conserver-8.1

Documentation Reference:

- For information on AIX installation, refer to [AIX 5L Version 5.2 Installation Guide and Reference \(SC23-4389-03\)](#)
- For information on CSM installation, refer to CSM for AIX5L Planning and Installation Guide
- For information on CSM configuration, refer to CSM for AIX5L Administration Guide
- [pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide for IBM eServer p5 servers](#)

For AIX5L Maintenance Package Download Information go to <http://www.ibm.com/servers/eserver/support/unixservers/aixfixes.html>

Notes:

- a. Check that the csm.gui.websm and csm.hpsnm filesets were installed as part of the AIX installation media and the option to "Automatically increase filesystems" was selected during the installation. These filesets are required for the operation of the High Performance Switch Network Manager (HPSNM) software on the CSM management server.
 - b. Check that the size of the /var directory on the CSM Management Server is approximately 1 Gigabyte to accommodate the HPSNM log files.
 - c. Additional information on installing the CSM software is also available in Chapter 6 ("Installation") of the [pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide for IBM eServer p5 servers \(SA38-0646-00\)](#).
- b. Apply RSCT and CSM Updates to CSM MS.
- Apply the RSCT and CSM PTF update images for AIX on the CSM Management Server.
 - RSCT PTF Download Information - refer to Component Update/Download Information for RSCT
 - CSM PTF Download Information - refer to Component Update/Download Information for

CSM

Level Check:

```
lslpp -L csm.client csm.core
csm.diagnostics csm.dsh
csm.gui.dcem
csm.gui.websm
csm.server
csm.hams
lslpp -L rsct.core.sensorm rsct.core.utils
rsct.exp.cimrm rsct.basic.rte
rsct.core.lprm rsct.core.rmc
rsct.core.sec rsct.basic.hacmp
rsct.core.auditrm
```

[Back to top](#)

Step 6: Define the nodes/LPARs in the cluster

To define the nodes/LPARS in the cluster, review the following:

- Refer to Chapter 6, "Software Installation", Set up the CSM cluster' in the [HPS Planning Guide](#)
- For detailed information about setting up the CSM cluster, refer to Chapter 8 of: [Cluster System Management \(CSM\) Library](#)
- CSM for AIX 5L Planning and Installation Guide

[Back to top](#)

Step 7: Assign IP Addresses for SNIs

If you are referencing "Step 5: Assign IP Addresses for SNIs" of the "Software Installation" section in Chapter 6 ("Installation") of the [pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide for IBM eServer p5 servers \(SA38-0646-00\)](#), the step cannot be implemented at this stage of the installation sequence due to a known problem with NIM whereby the defaults used by NIM for the number and size of technical large pages are not big enough to support the automatic configuration of more than one SNI as secondary adapters by NIM.

This step to assign IP Addresses to the SNIs will now be carried out after AIX is installed on the LPARs, the HPS/SNI updates are applied, and large page support enabled on the nodes.

[Back to top](#)

Step 8: Install AIX on all LPARs in the cluster

Refer to [pSeries HPS Planning, Installation and Service Guide](#) , Chapter 6, "Installation", Software Installation", "Install AIX on the nodes".

For the recommended service levels of AIX5L version 5.2 and AIX5L version 5.3 refer to the [Component Update/Download Information for AIX](#)

Note:

For the PE LP - DPCL is no longer a part of the IBM PE for AIX licensed program. Instead, DPCL is now available as an open source offering that supports PE. For more information and to download the DPCL open source project go to <http://oss.software.ibm.com/developerworks/opensource/dpcl>

This page should link to <http://sourceforge.net/projects/dpcl> where you can download the required code.

Because `/var` is a system data repository, system administrators should check periodically to maintain `/var` such that there is at least 30 Mb free [use `'df -k'`]. If `/var` is more than 75% full, use the command `du /var | sort -n` to find directories containing the most data.

Documentation Reference:

- For information on AIX installation, refer to [AIX 5L Version 5.2 Installation Guide and Reference \(SC23-4389-03\)](#)
- [pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide for IBM eServer p5 servers](#)
- Chapter 12, "Installing the AIX operating system on the nodes", in [CSM for AIX5L Planning and Installation Guide](#) which also contains details on installing the AIX operating system on the LPARs.

AIX5L Maintenance Package Download Information

To download the package, complete the following steps:

1. Go to <http://www.ibm.com/servers/eserver/support/unixservers/aixfixes.html>
2. Select Maintenance Packages and choose AIX52 or AIX53.
3. Specify your "Current level".
4. Specify "Desired level" and click "GO".

5. Follow instructions on this page.
6. Check AIX5L Critical Fixes.

Level check AIX by running the following command on the LPARs:

- o `oslevel -r`
- o `lspp -L bos.mp`

Refer to [Component Update/Download Information for AIX](#).

Note:

If filesets "csm.hpsnm" and "csm.server" are packaged with AIX5L, they should NOT be installed on the LPARs in the cluster.

[Back to top](#)

Step 9: Apply HPS/SNI Updates to all LPARs

Note that a filter installing AIX on the LPARs, only the HPS and SNI base level images are on LPARs in the cluster.

Download and apply the HPS/SNI PTF update referenced in the ["Detailed LP Check" for AIX5L 5.2 or AIX5L 5.3](#).

[Level Check:](#)

- o [lspp -Lc | egrep "HPS|sni" | cut -d : -f 2,3 | sed 's:/ /'](#)
- o [dsh "lspp -Lc | egrep "HPS|sni" | cut -d : -f 2,3 | sed 's:/ /' "|dshbak|more](#)

[Expected Levels are provided in Table 2](#) " HPS/SNI Detailed LP Level Check"

Notes:

- a. Notes on using TLP (Large Page) Settings in High Performance Computing (HPC) environment:
 - It is strongly recommended that users familiarize themselves with TLP basics and configuration options available.
 - The High Performance Switch adapter requires TLP usage and these TLP requirements are documented.
 - A formula to calculate the required TLP is outlined below.
 - You should also consult the section on the Large Page feature on AIX in

[AIX 5L Version 5.2 Performance Management Guide](#)

OR

[AIX 5L Version 5.3 Performance Management Guide](#)

- Users need to be aware of the usage of the LoadLevelerR pre-emption features with TLP (Large Pages) - specifically the fact that jobs that are using TLP that are pre-empted will essentially "lock up" the real memory the TLP's use, which is pinned by AIX. Unwise use of TLPs with LoadLeveler pre-emption can result in exhausting real memory available for jobs. If one wants LoadLeveler to schedule jobs based on the availability of large page, (especially if the job is going to run in mandatory Large Page mode) he/she may consider making use of the LoadLeveler consumable resource feature. The consumable resource function, which has already been around for several years, is documented in the LoadLeveler manual.

b. Notes on tuning Virtual Memory (VMM) Settings in an HPC environment:

Customers should be advised that the AIX VMM parameters (set by the vmo command) minfree and maxfree will most likely have to be adjusted (increased) in an HPC environment based on your cluster size, the amount of system memory, and number of processors per CEC. These settings when tuned properly will ensure enough memory remains available for core cluster infrastructure applications (RSCT, GPFS, LoadL). The recommended initial value for these tunables are minfree = 10000 and maxfree = 12000. Users are strongly urged to consult the following AIX documentation on virtual memory and vmstat tools and tune their system accordingly:

- http://publib16.boulder.ibm.com/doc_link/en_US/a_doc_lib/aixbman/prftungd/memperf.htm
- http://publib16.boulder.ibm.com/doc_link/en_US/a_doc_lib/aixbman/prftungd/memperf1.htm#i50853 The [AIX 5L Version 5.2 Performance Management Guide](#) should also be consulted.

Tuning these settings have been shown to help users avoid conditions where core cluster applications shut down and restart due to extensive blockage caused by "out of memory" issues. Keep in mind that all cluster applications should be designed and cluster tuned accordingly as to avoid oversubscribing to the real memory available.

After successful update of HPS/SNI filesets to the latest level, 64 bit kernel and technical large page support option must be enabled.

Set up your LPARs with 64-bit kernel

1. Check which kernel is currently in use: `bootinfo -K` a response of "32" is a 32bit Kernel
2. `ln -fs /usr/lib/boot/unix_64 /unix`
3. `ln -fs /usr/lib/boot/unix_64 /usr/lib/boot/unix`
4. Determine which rootvg hdisk contains the boot logical volume (usually hd5). This hdisk will be your "ipldevice".
 - `lspv |grep rootvg`

- hdisk0 009b982332a1f9b8 rootvg active
 - hdisk1 009b982332a2321a rootvg active
5. Issue: bosboot -ad /dev/<ipldevice> (eg. bosboot -ad /dev/hdisk0)
 6. Reboot: shutdown -Fr
 7. Verify 64 bit kernel is running after reboot: bootinfo -K 64

Set up Large Page Option

For configuration details see Large Page Support in [AIX 5L 5.2 Performance Management Guide](#) or

[AIX 5L 5.3 Performance Management Guide](#)

The number of TLP depends on customer configuration and relates to the number of windows required for each adapter(sni) plus any Large Pages used by other applications. Set up Large Page Option using the vmo command for each LPAR:

```
vmo -r -o v_pinshm=1 -o lgpg_size=16777216 -o lgpg_regions=YYY
```

where YYY is the amount of Technical Large Pages to export.

Note: you can use /usr/sni/aix5?/debugtools/sni_calc_lgpg_cnt to calculate the minimum number of TLP's.

Example:

On an LPAR with 2 sni's :

```
/usr/sni/aix5?/debugtools/sni_calc_lgpg_cnt
number_of_sni = 2
num_windows = 16
total_num_windows = 23
spool_size = 0x2000000
rpool_size = 0x2000000
Total required LPs == 0x39 [57]
```

You can also use sni_calc_lgpg_cnt to calculate TLP for a number of different vars.

Example:

```
/usr/sni/aix5?/debugtools/sni_calc_lgpg_cnt -a 8
number_of_sni = 8
num_windows = 16
total_num_windows = 23
```

```

spool_size = 0x2000000
rpool_size = 0x2000000
Total required LPs == 0xd5 [213]

```

Example:

To setup a node with 8 sni adapters:

```

16MB Large Page: lgpg_size = 16777216
256 Large Pages: lgpg_regions = 256
dsh <nodelist> "echo y|vmo -r -o v_pinshm=1 -o
lgpg_size=16777216 -o lgpg_regions=256"

```

(Use , echo y|vmo because otherwise vmo will prompt for verification to run bosboot). The following response will be generated:

```

Setting v_pinshm to 1 in nextboot file
Setting lgpg_size to 16777216 in nextboot file
Setting lgpg_regions to 256 in nextboot file
Warning: some changes will take effect only after a bosboot and a
reboot
Run bosboot now?
bosboot: Boot image is 19624 512 byte blocks.
Warning: changes will take effect only at next reboot

```

Note:

The vmtune sample program is being phased out and is not supported in future releases. It is replaced with the vmo command (for all the pure VMM parameters) and the ioo command (for all the I/O related parameters) which can be used to set most of the parameters that were previously set by vmtune. The -v flag has been added to vmstat to replace the -A flag which display counter values instead of tuning parameters. For AIX 5.2, a compatibility script calling vmo and ioo is provided to help the transition.

To Check that Large Page Option is set:

```

vmo -a|grep lg
lgpg_size = 16777216
lgpg_regions = YYY <where YYY is the amount of Technical Large Pages
to export
soft_min_lgpgs_vmpool = 0

```

Reboot LPARs. The LPARs should be rebooted as soon as possible to properly integrate the changes and to avoid disruption of current functionality.

Using memory by Remote Direct Memory Access (RDMA) and TLP

The following items outline information on using memory by RDMA and TLP:

- Real memory is divided into two categories - Small Pages and Large pages.
- It is the user's responsibility to achieve an optimal balance between the the two categories based on the expected and/or experienced needs of both SNI adapters memory requirements expressed in TLP and applications use of Small Pages as expressed in RDMA.
- TLP can allocate up to 75% of real memory.
- RDMA can pin and map up to 75% of small page application memory.
- Total Real Memory is a function of $N(\text{bytes of real mem}) = T(\text{bytes of real mem allocated to TLP}) + S(\text{bytes of real mem allocated to Small Pages})$.
- Small Page memory is a function of $S(\text{bytes of real mem allocated to Small Pages}) = N(\text{bytes of real mem}) - T(\text{bytes of real mem allocated to TLP})$.
- Large Page memory is a function of $T(\text{bytes of real mem allocated to TLP}) = N(\text{bytes of real mem}) - S(\text{bytes of real mem allocated to Small Pages})$.

The amount of small page memory can be calculated as follows:

- `lsattr -E -l sys0 -a realmem` returns the number of kbytes real memory, call this number A.
- `vmo -o lgpg_regions` returns the number of large pages, call this number B.
- Then $A * 1024 - B * 16 * 1024 * 1024$ is the amount of small page memory in bytes.

For example:

```
#!/bin/ksh
real_mem=`lsattr -E -l sys0 -a realmem|awk '{print $2}'`
lgpg_regions=`vmo -o lgpg_regions|awk '{print $3}'`
A=$(( real_mem * 1024 ))
B=$(( lgpg_regions * 16*1024*1024 ))
print "Real Mem=$A, TLP=$B; Small pages=$((A - B))"
```

- Real Mem=32212254720, TLP=4294967296; Small pages=27917287424
- The `rdma_xlat_limit` will limit the amount of memory that a user application can pin and map for use with RDMA.
- This pinning and mapping only survives as long as the job is executing. After it exits the memory is unpinned and freed.

Assign IP addresses to all adapter Switch Network Interfaces (SNIs)

Refer to [pSeries \(HPS\) Planning, Installation and Service Guide](#), Chapter 6, "Installation", "Assign IP addresses for SNIs" in the "Software Installation" section in of the for more details.

Note: The IP addresses and netmask are assigned to the SNIs using the *chdev* command.

Examples of the *chdev* command:

- `chdev -l sn0 -a netaddr=192.168.0.3 -a netmask=255.255.255.0 -a state=up`
- `chdev -l sn1 -a netaddr=192.168.1.3 -a netmask=255.255.255.0 -a state=up`
- `chdev -l ml0 -a netaddr=10.10.10.3 -a netmask=255.255.255.0 -a state=up`

[Back to top](#)

Step 10: Dynamic Host Configuration Protocol (DHCP)

Complete the following steps to for DHCP:

1. Disable Dynamic Host Configuration Protocol (DHCP) on the HMC and reboot the HMC.
2. Configure DHCP on CSM Management Server so that the lease time never expires. Details available in Chapter 7 ("Using Cluster-Ready Hardware Server (CRHS)") of the [CSM for AIX5L Administration Guide](#)
3. Configure Cluster Ready Hardware Server (CRHS) software on the CSM Management Server. Detailed instructions to configure CRHS on the CSM Management Server is available in Chapter 7 ("Using Cluster-Ready Hardware Server (CRHS)") of the [CSM for AIX5L Administration Guide](#).

Preliminary Notes:

The Secure Sockets Layer (SSL) RPM is a prerequisite for Secure Shell (SSH) which is required for the CRHS support.

Download Information:

- Download the SSL RPM from <http://www.ibm.com/link/oss.software.ibm.com/redirect.shtml/projects/opensshi>. The current version of SSL at the time of CSM 1.4.1 publication is `openssl-0.9.7d-2.aix5.1.ppc.rpm`.
- Download `openssh-3.8.1-p1_52.tar.gz` from <http://www.ibm.com/link/oss.software.ibm.com/redirect.shtml/projects/opensshi> and open the SSH tarball. The SSH tarball should include the following files:
 - `openssh.base.client 3.8.0.5200 COMMITTED Open Secure Shell Commands`
 - `openssh.base.server 3.8.0.5200 COMMITTED Open Secure Shell Server`
 - `openssh.license 3.8.0.5200 COMMITTED Open Secure Shell License`
 - `openssh.man.en_US 3.8.0.5200 COMMITTED Open Secure Shell`

- openssl.msg.en_US 3.8.0.5200 COMMITTED Open Secure Shell Messages
- openssl.base.client 3.8.0.5200 COMMITTED Open Secure Shell Commands
- openssl.base.server 3.8.0.5200 COMMITTED Open Secure Shell Server

Use the following documentation references:

- An overview of the CRHS setup is available in Chapter 4 ("System management components"), sub-section titled "Cluster-Ready Hardware Server concepts" of the [pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide for IBM eServer p5 servers](#) .
- Detailed instructions to configure CRHS on the CSM Management Server is available in Chapter 7 ("Using Cluster-Ready Hardware Server (CRHS)") of the [CSM for AIX5L Administration Guide](#).
- [CSM for AIX5L Planning and Installation Guide](#).
- [CSM download info \(select 'View'\)](#).

Level check

Verify the levels of RSCT and CSM on the CSM Management server:

- `lspp -L csm.client csm.core csm.diagnostics
csm.dsh csm.gui.dcem csm.gui.websm
csm.server csm.hams`
 - [AIX52 / AIX53 CSM Level check](#)
- `lspp -L rsct.core.sensorrm rsct.core.utils
rsct.exp.cimrm rsct.basic.rte
rsct.core.lprm rsct.core.rmc
rsct.core.sec rsct.basic.hacmp
rsct.core.auditrm`
 - [AIX52 RSCT level check](#)
 - [AIX53 RSCT level check](#)

Notes:

- A. Please see the following CSM webpage for packages and known issues
<http://www14.software.ibm.com/webapp/set2/sas/f/csm/download/home.html>
- B. The CSM Management Server should have the right levels of RSCT and CSM PTFs after in order to configure CRHS.
- C. From the [CSM FAQs](#):
 - How should the service network be configured on the HMC as part of the Cluster Ready Hardware Server setup after the DHCP server migration to the CSM Management Server?
 - The service network interface on the HMC must be configured as "static" to prevent the

DNS nameserver in the /etc/resolv.conf file from getting modified during an HMC update. Use the following steps on the HMC to enable a static service network on the HMC:

1. On the HMC GUI, in the "HMC Management" panel, select the "HMC Configuration" option.
2. Select the "Customize Network Settings" option.
3. Under the "LAN Adapters" tab, select the service network interface for your HMC, example, "Ethernet eth0" and click on "Details"
4. Under the "LAN Adapter" tab, in the "Local Area Network Information sub-section, select the "open" option.
5. In the "DHCP Client/IP Address" sub-section, select the "Specify an IP address" option and enter the TCP/IP interface address and network mask of the service network.
6. Click *OK*.
7. The HMC will have to be rebooted in order for the changes to be in effect.

[Back to top](#)

Step 11: Power off CECs

Ensure all CECs in the cluster are in the powered off state before proceeding.

[Back to top](#)

Step 12: Configure the High Performance Switch Network Manager (HPSNM)

Configure the High Performance Switch Network Manager (HPSNM) software, bring the network online and report installation complete.

1. Install the HPSNM_Serv_Pack.zip on the HMC and apply the csm.hpsnm PTF fileset to the CSM MS.
2. Bring the network online and report installation complete using the instructions in the "Bringing the network online and reporting Installation Complete" section, in Chapter 6 ("Installation") of the [pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide for IBM eServer p5 servers](#).

Download Information:

For the HMC: Install the HPSNM Service Pack package from the HPSNM/IBMNM fixes tab of the appropriate HMC Level as indicated by the System Firmware shown in the "[Component Update/Download Information Section](#)" for HPSNM:

- [235](#)
- [240](#)

For the CSM MS:

Download and Install csm.hpsnm AIX PTF as indicated by the "[Component Update/Download Information Section](#)" in the HPSNM section.

CSM Packages for AIX and Linux are available at:

<https://www14.software.ibm.com/webapp/set2/sas/f/csm/download/home.html>

Installation on the HMC:

- The HPSNM_Serv_Pack_<n>.zip is installed on all HMCs in the cluster, and each HMC should be rebooted.
 - Install the zip file on all HMCs from the HMC GUI using the following options:
 - Licensed Internal Code Maintenance -> HMC Code Update -> Install Corrective Service
 - Select the appropriate install method:
 - Apply corrective service from removable media if HPSNM_Serv_Pack_2.zip was downloaded to a removable media
- OR
- download the corrective service file from a remote system, and then apply the downloaded service file if the zip file was downloaded to a remote system.
 - Reboot all HMCs with the HPSNM Service Pack installed.

Installation on the CSM Management Server (CSM MS):

- The csm.hpsnm base fileset is installed on the CSM Management Server as part of the CSM MS installation with AIX5L Maintenance package (Use the "Automatically increase filesystems" option when doing this installation).
- Apply the csm.hpsnm PTF fileset on the CSM MS.

Note: In order to accommodate the HPSNM log files, /var on the CSM Management Server should be increased to 1 Gigabyte.

More Configuration:

- ELA Master Selection. This is needed for tracelogd.
- /opt/csm/hpsnm/bin/chswelamast -q # Used to query the ELA Master and Backup HMCs
- /opt/csm/hpsnm/bin/chswelamast -m <hostname or IP address> [-b <hostname or IP address>] # Used to set the ELA Master and Backup HMCs.

Switch Topology Selection:

- `/opt/csm/hpsnm/bin/chswlogtop -n1 <n1's_topology> -n2 <n2's_topology> # This should be done for each network.`
- `/opt/csm/hpsnm/bin/chswlogtop -h` Shows syntax, including supported topologies.
- `/opt/csm/hpsnm/bin/chswlogtop -n <network_number> none` To clear a network topology.

Activation and Verification:

- Verify that all switch frames are powered.
- Verify that all CECs are powered off.
- Complete the procedure to bring the network online and report installation complete, otherwise known as, Service code 20 in the US. Refer to the section titled, "Bringing the network online and reporting Installation Complete" in Chapter 6 ("Installation") of the [pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide for IBM eServer p5 servers \(SA38-0646-00\)](#) for complete details.
- Examine switch link status by running the following command: `/opt/csm/hpsnm/bin/lsswtopol {-n network}`
- Examine endpoint link status by running the following command: `/opt/csm/hpsnm/bin/lsswendpt`

Looking for Events:

On the ELA Master HMC GUI, select the following options to view the reported events/errors: Service Applications --> Service Focal Point --> Manage Serviceable Events

Collecting Data:

In case you need to report a problem, run `/opt/csm/hpsnm/bin/fnm.snap` to generate a snap. Look for core dumps in `/var/opt/csm/hpsnm/log/`

Documentation Reference:

- [pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide for IBM eServer p5 servers](#) .

Level Check:

On the CSM MS, do either one of the following:

- Use `lspp -L csm.hpsnm` to verify csm.hpsnm successfully installed. See CSM in the HPSNM component

- Create a service login ID on the CSM Management Server.

For instructions on creating a service login ID, refer to "Step 21: Create hardware service representative login ID" in Chapter 5 ("Installing the management server on AIX") of the [CSM for AIX5L V1.4.1 Planning and Installation Guide](#).

Note:

The service login ID, with the appropriate permissions, is necessary for the IBM Service Representative to access service-related HPSNM functions on the CSM Management Server without requiring root authority.

[Back to top](#)

Health check/planning help

- [FLRT: Fix Level Recommendation Tool](#)
- [IBM Systems compatibility information](#)

- [Cluster support](#)
- [High Performance Switch](#)

IBM HPS POWER5 Readme - Service Pack 17

AIX LP

Anchor link navigation

- [Introduction](#)
- [Supporting documentation](#)
- [Component updates/download information](#)
- [Upgrade / Installation sequences](#)

- **[AIX LP](#)**
- [Known problems and workarounds](#)
- [Readme archive](#)

AIX Licensed Program (LP)

The LP fileset fix levels listed in the [Detailed LP Level Check](#) table below indicate the latest available levels at the time this HPS Service Pack was released. The published LP fix levels in the table below are intended to be the minimum levels for this Service Pack, unless explicitly stated otherwise. They are available from [IBM Fix Central](#) and [Power Systems cluster software support fixes](#).

To order the latest filesets for all LP components of this HPS Service Pack, click the APAR link for your Version of AIX:

AIX 5.3: [APAR IZ22887](#)

AIX 5.2: [APAR IZ22887](#)

As a convenience, the table contains an "order the latest" APAR number for each of the LPs.

Clicking on the LP Component link will open a new page in FixCentral. There you will be presented with an opportunity to get the latest available fixes for that component.

Please be aware of the "Replace superseded fixes with the latest" packaging option. If you choose to get the exact levels as stated in this Service Pack, you must uncheck the option before proceeding to download.

The selected APAR package will contain fix levels that are the same or higher than the levels listed below.

To check the LP service levels, on each logical partition issue:

```
lspp -Lc | egrep "vsd|LAPI|HPS|sni|ppe|LoadL|mmfs|rsct|csm|essl|pessl" | cut -d : -f 2,3 | sed 's:/ /'
```

LP	AIX 5.3	AIX 5.2
VSD	4.1.0	4.1.0
LAPI	2.4.5	2.3.3
HPS/SNI (Devices)	1.2.0	1.1.3
PE	4.3.2	4.2.2
LoadLeveler	3.4.3	3.3.2
GPFS	3.1.0	2.3.0
	3.2.1	
CSM	1.6.0	1.5.1
	1.7.0	1.6.0
RSCT	2.4.9	2.3.12
ESSL	4.2.0	4.2.0
	4.3.0	4.3.0
Parallel ESSL	3.3.0	3.2.0

Detailed LP Level Check

Depending on the LP's that are installed, your output should match the LP fileset levels listed in the Detailed LP Level Check table. Use the following Detailed LP Level Check table to order the latest APARs that are available from [FixCentral](#).

LP	Release	Component ID	APAR	Level check	Info
VSD	4.1.0	5765G2602	IZ22869	rsct.vsd.cmds 4.1.0.18 rsct.vsd.rvsd 4.1.0.19 rsct.vsd.vsdd 4.1.0.23 rsct.vsd.vsdrm 4.1.0.7	View
LAPI	2.3.3	5765G2601	IZ08270	rsct.lapi.nam 2.3.3.3 rsct.lapi.rte 2.3.3.10 rsct.lapi.samp 2.3.3.0	View
AIX52 only					

LAPI	2.4.5	5765G2601	IZ22871	rsct.lapi.nam 2.4.4.0 rsct.lapi.rte 2.4.5.1 rsct.lapi.samp 2.4.5.0	View
AIX53 only					
HPS/SNI (Devices)	1.1.3	5765G2400	IZ22872	devices.chrp.IBM.HPS. hpsfu 1.1.3.0 devices.chrp.IBM.HPS.rte 1.1.3.14 devices.common.IBM.sni.ml 1.1.3.2 devices.common.IBM.sni.ntbl 1.1.3.1 devices.common.IBM.sni.rte 1.1.3.12 devices.msg.en_US.chrp.IBM. HPS.hpsfu 1.1.3.0 devices.msg.en_US.chrp.IBM. HPS.rte 1.1.3.1 devices.msg.en_US.common.IBM. sni.ml 1.1.3.1 devices.msg.en_US.common.IBM. sni.ntbl 1.1.3.1 devices.msg.en_US.common.IBM. sni.rte 1.1.3.0	View
AIX52 only					
HPS/SNI (Devices)	1.2.0	5765G2400	IZ22873	devices.chrp.IBM.HPS. rte 1.2.0.12 devices.common.IBM.sni. ml 1.2.0.5 devices.common.IBM.sni. ntbl 1.2.0.0 devices.common.IBM.sni. rte 1.2.0.11 devices.msg.en_US.chrp.IBM. HPS.rte 1.2.0.2 devices.msg.en_US.common.IBM. sni.ml 1.2.0.1 devices.msg.en_US.common.IBM. sni.ntbl 1.2.0.1	View
AIX53 only					
PE	4.2.2	5765F8300	IZ22874	ppe.man 4.2.2.2 ppe.perf 4.2.2.3 ppe.poe 4.2.2.15 ppe.pvt 4.2.2.0	View
PE	4.3.2	5765F8300	IZ22875	ppe.man 4.3.2.0 ppe.perf 4.3.2.0 ppe.poe 4.3.2.1 ppe.pvt 4.3.2.0	View
AIX53 only					

LoadLeveler	3.3.2	5765E6900	IZ22876	LoadL.full 3.3.2.16 LoadL.msg.En_US 3.3.2.5 LoadL.msg.en_US 3.3.2.5 LoadL.so 3.3.2.16	View
LoadLeveler	3.4.2	5765E6900	IZ22877	LoadL.full 3.4.3.1 LoadL.msg.en_US 3.4.3.1 LoadL.so 3.4.2.1 LoadL.webui 3.4.3.0 LoadL.tguides 3.4.3.0	View
AIX53 only					
GPFS	2.3.0	5765F64AP	IZ22878	gpfs.base 2.3.0.30 gpfs.docs.data 2.3.0.8 gpfs.msg.en_US 2.3.0.14	View
GPFS	3.1.0	5765G66P	IZ22880	gpfs.base 3.1.0.20 gpfs.docs.data 3.1.0.5 gpfs.msg.en_US 3.1.0.12	View
GPFS	3.2.0	5765G66P	IZ22881	gpfs.base 3.2.1.2 gpfs.docs.data 3.2.0.1 gpfs.msg.en_US 3.2.1.2	View
CSM	1.5.1	5765F6700	IZ22882	csm.bluegene 1.5.1.1 csm.client 1.5.1.4 csm.core 1.5.1.4 csm.deploy 1.5.1.4 csm.diagnostics 1.5.1.0 csm.dsh 1.5.1.7 csm.essl 1.5.1.0 csm.gpfs 1.5.1.0 csm.gui.dcem 1.5.1.0 csm.gui.websm 1.5.1.0 csm.hams 1.5.1.1 csm.ll 1.5.1.0 csm.pe 1.5.1.0 csm.pessl 1.5.1.0 csm.server 1.5.1.7 (Install only on CSM Management Server) csm.hpsnm 1.4.1.18	View

RPMS:

expect 5.32-1
openCIMOM 0.8-1
openssl 0.9.7g-1
AIX-rpm 5.2.0.95-2
tcl 8.3.3-1
tk 8.3.3-1

CSM	1.6.0	5765F6700	IZ22883	<pre> conserver 8.1.7-2 csm.pessl 1.6.0.11 csm.pe 1.6.0.10 csm.ll 1.6.0.10 csm.hc_utils 1.6.0.11 csm.hams 1.6.0.10 csm.gui.websm 1.6.0.10 csm.gui.dcem 1.6.0.10 csm.gpfs 1.6.0.10 csm.essl 1.6.0.11 csm.dsh 1.6.0.13 csm.diagnostics 1.6.0.10 csm.deploy 1.6.0.11 csm.core 1.6.0.12 csm.client 1.6.0.13 csm.bluegene 1.6.0.10 csm.server 1.6.0.16 (install only on CSM Management Servers) </pre>	View
AIX52 only					
				<pre> RPMS: expect 5.32-1 openCIMOM 0.8-1 openssl 0.9.7g-1 AIX-rpm 5.2.0.95-1 tcl 8.3.3-1 tk 8.3.3-1 conserver 8.1.7-2 </pre>	
CSM	1.7.0	5765F6700	IZ22900	<pre> csm.bluegene 1.7.0.11 csm.client 1.7.0.12 csm.core 1.7.0.12 csm.deploy 1.7.0.11 csm.diagnostics 1.7.0.11 csm.dsh 1.7.0.12 csm.essl 1.7.0.11 csm.gpfs 1.7.0.11 csm.gui.dcem 1.7.0.11 csm.gui.websm 1.7.0.11 csm.hams 1.7.0.11 csm.hc_utils 1.7.0.11 csm.ll 1.7.0.11 csm.pe 1.7.0.11 </pre>	View
AIX53 only					

csm.pessl 1.7.0.11
 csm.hpsnm 1.4.1.18
 csm.server 1.7.0.12 (install only
 on CSM Management Servers)

RPMS:

expect 5.32-1
 openCIMOM 0.8-1
 openssl 0.9.7d-2
 AIX-rpm-5.3.8.0-9
 tcl 8.3.3-1
 tk 8.4.7-3
 conserver 8.1.7-2

[RSCT](#) 2.3.12 5765F07AP IZ22885 devices.chrp.base.
 ServiceRM 1.2.0.117
 rsct.basic.hacmp 2.3.12.0
 rsct.basic.rte 2.3.12.3
 rsct.basic.sp 2.3.12.0
 rsct.compat.basic.
 hacmp 2.3.12.0
 rsct.compat.basic.
 rte 2.3.12.1
 rsct.compat.basic.
 sp 2.3.12.0
 rsct.compat.clients.
 hacmp 2.3.12.0
 rsct.compat.clients.
 rte 2.3.12.0
 rsct.compat.clients.
 sp 2.3.12.0
 rsct.core.auditrm 2.3.12.2
 rsct.core.errm 2.3.12.2
 rsct.core.fsrms 2.3.12.1
 rsct.core.gui 2.3.12.0
 rsct.core.hostrm 2.3.12.1
 rsct.core.lprm 2.3.12.1
 rsct.core.rmc 2.3.12.2
 rsct.core.sec 2.3.12.1
 rsct.core.sensorrm 2.3.12.1
 rsct.core.sr 2.3.12.1
 rsct.core.utils 2.3.12.3
 rsct.opt.saf.amf 2.3.12.0

[View](#)

rsct.opt.storagerm 2.3.12.2

[RSCT](#) 2.4.9 5765F07AP IZ22886
AIX53 only

devices.chrp.base.ServiceRM 1.3.0.80
rsct.basic.hacmp 2.4.9.1
rsct.basic.rte 2.4.9.1
rsct.basic.sp 2.4.9.1
rsct.compat.basic.hacmp 2.4.9.0
rsct.compat.basic.rte 2.4.9.0
rsct.compat.basic.sp 2.4.9.0
rsct.compat.clients.hacmp 2.4.9.0
rsct.compat.clients.rte 2.4.9.0
rsct.compat.clients.sp 2.4.9.0
rsct.core.auditrm 2.4.9.0
rsct.core.errm 2.4.9.0
rsct.core.fsrn 2.4.9.0
rsct.core.gui 2.4.9.0
rsct.core.hostrm 2.4.9.1
rsct.core.lprm 2.4.9.0
rsct.core.rmc 2.4.9.1
rsct.core.sec 2.4.9.1
rsct.core.sensorrm 2.4.9.1
rsct.core.sr 2.4.9.0
rsct.core.utils 2.4.9.1
rsct.opt.storagerm 2.4.9.0
rsct.opt.fencerm.hmc 2.4.9.1

[View](#)

[Parallel](#) 3.2.0 5765F8400 PK21664
[ESSL](#)

pessl.rte.common 3.2.0.1
pessl.rte.hv 3.2.0.1
pessl.rte.rs1 3.2.0.1
pessl.rte.smp 3.2.0.1
pessl.man.en_US 3.2.0.1
pessl.rte.mp 3.2.0.0
pessl.rte.rs2 3.2.0.0
pessl.rte.up 3.2.0.0
pessl.loc.license 3.2.0.0

[View](#)

Parallel	3.3.0	5765F8400	PK52541	<p> pessl.rte.common 3.3.0.2 pessl.rte.hv 3.3.0.2 pessl.rte.rs1 3.3.0.2 pessl.rte.smp 3.3.0.2 pessl.man.en_US 3.3.0.0 pessl.rte.mp 3.3.0.0 pessl.rte.rs2 3.3.0.0 pessl.rte.up 3.3.0.0 pessl.loc.license 3.3.0.0 </p>	View
ESSL	420	5765F8200		<p> essl.rte.common 4.2.0.5 essl.rte.rs1 4.2.0.5 AIX 5.3 essl.rte.rs2 4.2.0.0 PK20638essl.rte.smp 4.2.0.5 PK19344essl.rte.mp 4.2.0.0 AIX 5.2 essl.rte.up 4.2.0.0 PK20638essl.man.en_US 4.2.0.0 PK19344essl.loc.license 4.2.0.0 essl.msg.en_US 4.2.0.1 essl.msg.En_US 4.2.0.1 essl.msg.en_US 4.3.0.0 AIX 5.3 essl.msg.En_US 4.3.0.0 PK52484essl.rte.common 4.3.0.2 AIX 5.2 essl.rte.rs1 4.3.0.2 PK52484essl.rte.smp 4.3.0.2 </p>	View
ESSL	430	PK52484		<p> AIX 5.3 essl.msg.en_US 4.3.0.0 essl.msg.En_US 4.3.0.0 PK52484essl.rte.common 4.3.0.2 AIX 5.2 essl.rte.rs1 4.3.0.2 PK52484essl.rte.smp 4.3.0.2 </p>	View

[Back to top](#)

Health check/planning help

- [FLRT: Fix Level Recommendation Tool](#)
- [IBM Systems compatibility information](#)

- [Cluster support](#)
- [High Performance Switch](#)

IBM HPS POWER5 Readme - Service Pack 17

Known problems and workarounds

Anchor link navigation

- [Introduction](#)
- [Supporting documentation](#)
- [Component updates/download information](#)
- [Upgrade / Installation sequences](#)

- [AIX LP](#)
- **[Known problems and workarounds](#)**
- [Readme archive](#)

Restrictions/Known Problems/Workarounds

This pSeries HPS POWER5 release does not support GA5 (230), or GA6 (235) concurrent firmware installation in an operational HPS POWER5 cluster pending further testing of this function at these levels. However, this pSeries HPS POWER5 release does support GA7 (240) concurrent firmware installation in an operational HPS POWER5 cluster.

The use of High Availability Management Server is restricted with this release pending further testing.

Customers who upgrade to this pSeries HPS POWER5 Service Pack and try to enable the redundant FSP feature with invalid Klingon configurations may result in ambiguous to troublesome errors.

Upon application of csm.hpsnm APAR IY88538 (available as of csm.hpsnm 1.4.1.16, included in p5 HPS SP11), customers can safely reenble the High Performance Switch (HPS) Time Of Day (TOD) support. In a previous service bulletin/flash, HPSNM Development recommended that customers disable this support by adding the following to "/etc/environment "on the Management Station:
"HPSNM_TOD_SUPPORT=off".

Customers that wish to reenble HPS TOD support should apply IY88538, remove the environment variable from "/etc/environment" and restart HPSNM. For customers that have not disabled the HPS

TOD support, no action is required following the application of IY88538 , but be aware there will be a momentary interruption of the TOD on the first restart (only) of HPSNM. This could impact the POE co-scheduler, as well as applications that use the MPI _WTIME subroutine that reads the TOD.

Retain Tip Abstracts

1. [HPS/SNI Protocol stack requires all sni in all servers to be on the same network](#)
2. [When updating Licensed Internal Code, if the "High Performance Switches"](#)
3. [No support in Service Focal Point for collecting VPD for High Performance Switches](#)
4. [B1818600 error is posted from Perc](#)
5. [Focal Point \(SFP\) flags a switch board or riser with a bad BIST and/or bad Signature.](#)
6. [The HPS Network Manager End-Point View and lsswendpt command show cage 16 as cage 0.](#)
7. [VPD databases lost on the HMC](#)
8. [When verifying switch chip ports, you should look for Down:No Signal.](#)
9. [Diagnostic SRC D103B0FF in error logs](#)
10. [Cluster Ready Hardware Server data may not be restored after an upgrade install of the HMC](#)

Retain Tip Abstract:

HPS/SNI Protocol stack requires all sni of same number in all servers to be on the same network in dual-network HPS clusters.

Search Keys:

Miswire, Miscable, SNI, ping problem, intermittent packet, intermittent ping, packet loss/

Problem Description:

If the corresponding SNIs in the servers are not on the same network, packets over the switch will intermittently fail. For example, if sni0 in server A is on network 1 and sni0 in server B is on network 2, you will have packets intermittently fail. There is no indication of this in the HPSNM GUI. Furthermore, if an SNI fails in such a way that it is not recognized by AIX, and it would not be the highest numbered sni in the configuration, all SNIs that would have followed it in the tree are now one less than they would have been previously. This causes the appearance of a miswire.

Customer Impact:

Performance degradation and potential application failure.

Workaround:

While performing the "Bringing the Network online" procedure in the HPS Planning, Install and Service Guide, the following steps should be performed just before running the host verification tools: If intermittent packet failure is observed after installation, use the following steps to:

A. Check for incorrectly cabled SNIs:

1. Go to the CSM MS and check the `"/var/opt/csm/hpsnm/log/topology.map"` file for: "There are X inconsistently cabled adapters". They are listed in HPSNM logical locations and indicate frame, cage and chip locations.
2. Record the locations of the inconsistently cabled adapters.
3. Open the HPSNM Endpoint View on the CSM MS GUI.
4. Look up the SNIs frame, cage and chip in the HPSNM and cross-reference to the physical location on the GUI panel.
5. Select this SNI link.
6. Click "Selected-Properties" on the menu
7. Note the network for this SNI.
8. Check the cabling instructions and recable as instructed:
 - If the cable instructions indicate that this SNI is cabled correctly, the cable planning should be redone such that this SNI is placed on the same network as the corresponding SNIs in the other servers.
 - Keep in mind that it may be possible that the entire switch has been wired into the incorrect network. In such a case, all SNIs connected to this switch will be considered to be cabled inconsistently.
9. After making changes, you should check GUI status and SFP for errors that may have been caused by recabling.

B. Check for unconfigured SNIs:

1. If the "Bringing the Network Online" procedures indicated any SNIs that were not configured and visible by AIX, perform the following on those SNIs:
2. For each server that had unconfigured SNIs:
 1. `mv /usr/sni/aix52/cfgsni /usr/sni/aix52/cfgsni.orig`
 2. Reboot server
 3. For each snix on the server: `rmdev -d -l snix`
 4. `mv /usr/sni/aix52/cfgsni.orig /usr/sni/aix52/cfgsni`
 5. `cfgmgr -s`

C. To double-check for unconfigured SNIs causing others to be on the incorrect netid.

(assumption: all discrepancies found in the topology.map were fixed)

1. Run the following command string on the CSM MS. This assumes that all LPARs are accessible via "dsh -av"
(Be mindful of the placement of quotation marks).
for x = 0 1 2 3 4 5 6 7; do echo sni\$x;
dsh -av `"/usr/sni/aix52/debugtools/sni_qry -l sni$x | grep netid" 2>/dev/null | uniq -f 3;`
done
2. There is a problem if more than one adapter is listed under each "sniX" heading. If this is the case, specific queries for the SNIs that may potentially be bad should be carried out.

3. With a knowledge of the correct network for sniX, run the following command on the CSM MS:
(be mindful of quotation marks and backslashes)
`dsh -av '/usr/sni/aix52/debugtools/sni_qry sniX | grep "netid: [wrong network]"'`
where you give the wrong network as the [wrong network] part of the grep.
For example, if sni0 should be network 1:
`dsh -av "/usr/sni/aix52/debugtools/sni_qry sni0 | grep "netid: 2""`
4. Now that you have the list of LPARs with problem adapters, you can recover on each LPAR:
 - a. `mv /usr/sni/aix52/cfgsni /usr/sni/aix52/cfgsni.orig`
 - b. Reboot server.
 - c. For each snix on the server: `rmdev -d -l sniX`
 - d. `mv /usr/sni/aix52/cfgsni.orig /usr/sni/aix52/cfgsni`
 - e. `cfgmgr -s`

[Back to top](#)

Retain Tip Abstract:

When updating Licensed Internal Code, if the "High Performance Switches" target is selected with one or more managed system targets, a `java.lang.ArrayStoreException` occurs immediately.

Problem Description:

Licensed Internal Code update fails with a `java.lang.ArrayStoreException` when the "High Performance Switches" target is selected with one or more managed system target. A serviceable event with refcode E302F9D2 will be recorded in Service Focal Point. Licensed Internal Code is not updated on any of the specified targets.

Customer Impact:

Licensed Internal Code cannot be updated simultaneously on managed system and "High Performance Switches" targets.

Workaround:

Update Licensed Internal Code for the "High Performances Switches" target separately from any other targets. Updating multiple managed system targets simultaneously continues to work correctly.

[Back to top](#)

Retain Tip Abstract:

No support in Service Focal Point for collecting VPD for High Performance Switches

Problem Description:

The Collect VPD Information function under the Service Focal Point component of the Hardware Management Console does not support the collection of vital product data for High Performance Switch clusters.

Customer Impact:

The customer will have to manually execute a command on the HMC to collect VPD for High Performance Switch cluster.

Workaround:

Execute the following command on the HMC:

```
/opt/hsc/bin/vpdfs -x cluster
```

The command will produce an xml file in the working directory named *7045-SW4-bycluster.xml*. This file contains the VPD for the switch cluster.

[Back to top](#)

Retain Tip Abstract:

B1818600 error is posted from Perc.

Problem Description:

A B1818600 error is posted from Perc with a user section showing:

```
Process name: netsSlp
```

Customer Impact:

There is no impact to the customer or the equipment.

Workaround: No workaround is required since this is a recoverable situation.

[Back to top](#)

Retain Tip Abstract:

Service Focal Point (SFP) flags a switch board or riser with a bad BIST and/or bad Signature.

Search Keys: *Bad BIST, Bad Signature.*

Problem Description:

A bad BIST error for a switch board or a riser is logged by fnmd when the self test computes a value that does not match what fnmd is expecting. This may or may not lead to subsequent errors.

Customer Impact:

Most likely no impact. Wait for other error types before taking corrective action. Return FNM_Init.log or an fnmd snap to IBM for analysis.

[Back to top](#)

Retain Tip Abstract:

The HPS Network Manager End-Point View and lsswendpt command show cage 16 as cage 0.

Problem Description:

If a frame has adapters in cage 16, the HPS Network Manager End-Point

View and lsswendpt command will show these adapters in cage 0.

Customer Impact:

Customer is presented incorrect information, but there is no loss of function.

Workaround:

None

[Back to top](#)

Retain Tip Abstract:

VPD databases lost on the HMC Search Keys:

Problem Description:

This is a problem which FVT and MPV have both hit. What happens is the following:

1. Tracelogd on the CSM and on the HMC are both running.
2. Tracelogd on the CSM starts sending the VPD files and manages to get one or two across.
3. Tracelogd on the HMC dies and starts to come back up. (it takes 5 seconds before tracelogd fully comes up and re-connects with hardware server and can start receiving messages again)
4. Tracelogd on the CSM realizes that there's a problem sending the files and goes into an error loop.

Unfortunately, this error loop will time out before the 5 seconds are up and the rest of the VPD files will not get across. This could possibly leave either a database completely off the HMC or leave only half of one there.

Customer Impact:

A database could either only be copied halfway onto the HMC or left off completely. This means that ELA could possibly not have the MTMS information needed to let a IBM SSR know where a problem is located.

Workaround:

The workaround is to take a look at the databases on the HMC. If one of the three databases (sma_vpd.db, chassis_vpd.db and bpa_vpd.db) is missing, copy it over from the CSM (you can use scp or ftp for example). CSM directory: /var/opt/csm/hpsnm/data HMC directory: /opt/hsc/data

[Back to top](#)

Retain Tip Abstract:

When verifying switch chip ports, you should look for Down:No Signal. Search Keys: No Signal, Down: No Signal

Problem Description:

Down:No signal indicates that a clock or a light-present signal is not seen on a port. This can mean that a switch link has no cable, or that the link is faulty. There is a missing check for this in the HPS Service Doc procedure, Bringing the Network Online, Step 3: Verify the network, Task 1: verify SNI and switch

chip ports, 1) verify switch ports, step 2.b

Customer Impact:

If the link is faulty, it will be un-useable and may degrade performance.

Workaround:

When executing the HPS Service Doc procedure: Bringing the Network Online

Step 3: Verify the network

Task 1: verify SNI and switch chip ports

1) verify switch ports, step 2.b

you should also filter on "No Signal". If you find one of these, go to the "Network Status Codes on HPSNM" table in Appendix A, and follow the instructions for "Down:No Signal".

[Back to top](#)

Retain Tip Abstract:

DiagnosticSRC D103B0FF in error logs

Search Keys:

D103BOFF Predictive Error

Problem Description:

An informational error was incorrectly categorized as a predictive failure.

Customer Impact:

Unnecessary call for service action.

Workaround:

Ignore the error. It is not a true software problem. The indication is that a program attempted to get information about the state of the hardware when the hardware was powered off.

[Back to top](#)

Retain Tip Abstract:

Cluster Ready Hardware Server data may not be restored after an upgrade install of the HMC resulting

in loss of CEC and switch configuration on the HMC and from the CSM management server.

Search Keys:

Cluster Ready Hardware Server (CRHS), HPS, HMC

Systems Affected:

CECs, HMCs HPS switch configured to use Cluster Ready Hardware Server

Problem Description:

It has been observed in several instances of upgrading the HMC from 01SF230_xxx_120 to a later SP level or to 01SF235_xxx_160 that the RSCT data which contains the Cluster Ready Hardware Server configuration is not restored. This happened when the HMC system clock was adjusted as specified in the update process.

Customer Impact:

HMC and CEC Configuration may be lost. Customers will not see CEC configuration in the GUI although the CEC/LPARS should still be functioning prior to the install.

Workaround:

1. If root access is enabled to the HMC then the data can still be restored by the following:
 - a. log in as root
 - b. check if the directory `"/var/ct.backup"` exists with a timestamp of when the save data was performed.
 - c. If it exists then - run the command `"/usr/sbin/rsct/bin/ctrestore"`
 - d. wait a few minutes for the ClusterPeerDomain to be restored and activated.
 - e. verify that the Cluster Ready Hardware Server data has been restored: run the command `"lsrsrc IBM.hw_manager"`

The output should show the HMCs in the cluster and the CSM_MS object.

If the directory does not exist then:

- a. `mount /mnt/upgrade`
- b. verify the `/mnt/upgrade/RSCTSaveUpgrade.tar` file exists.
- c. If the tar file exists then run:
`tar -xvf /mnt/upgrade/RSCTSaveUpgrade.tar`
- d. check if the directory `"/var/ct.backup"` exists.
- e. If it exists then run the command `"/usr/sbin/rsct/bin/ctrestore"`
- f. Wait a few minutes for the ClusterPeerDomain to be restored and activated.
- g. Verify that the Cluster Ready Hardware Server data has been restored by running the command `"lsrsrc IBM.hw_manager"`

The output should show the HMCs in the cluster and the CSM_MS object.

2. If the above was not successful or if the customer does not have root access to the HMC then the Cluster Ready Hardware Server data must be recreated as it originally was using the addpeer command from the CSM management server.

Plan to Resolve:

Plan to fix in a subsequent Service Pack.

[Back to top](#)

Health check/planning help

- [FLRT: Fix Level Recommendation Tool](#)
- [IBM Systems compatibility information](#)

- [Cluster support](#)
- [High Performance Switch](#)

IBM HPS POWER5 Readme - Service Pack 17

Archive

Anchor link navigation

- [Introduction](#)
- [Supporting documentation](#)
- [Component updates/download information](#)
- [Upgrade / Installation sequences](#)

- [AIX LP](#)
- [Known problems and workarounds](#)
- **[Readme archive](#)**

Available archives

The following links are provided for archival and historical purposes only. Please refer to the latest service pack document for interactive capabilities and available software.

[IBM HPS POWER5 FLASH/Readme - Service Pack 16](#)

[IBM HPS POWER5 FLASH/Readme - Service Pack 15 Addendum 1](#)

[IBM HPS POWER5 FLASH/Readme - Service Pack 15](#)

[IBM HPS POWER5 FLASH/Readme - Service Pack 14 Addendum 1](#)

[IBM HPS POWER5 FLASH/Readme - Service Pack 14](#)

[IBM HPS POWER5 FLASH/Readme - Service Pack 13](#)

[IBM HPS POWER5 FLASH/Readme - Service Pack 12](#)

[IBM HPS POWER5 FLASH/Readme - Service Pack 11](#)

[IBM HPS POWER5 FLASH/Readme - Service Pack 10](#)

[IBM HPS POWER5 FLASH/Readme - Service Pack 9](#)

[IBM HPS POWER5 FLASH/Readme - Service Pack 8](#)

[IBM HPS POWER5 FLASH/Readme - Service Pack 7](#)

[IBM HPS POWER5 FLASH/Readme - Service Pack 6](#)

[IBM HPS POWER5 FLASH/Readme - Service Pack 5](#)

[IBM HPS POWER5 FLASH/Readme - Service Pack 4 \(see SP5\)](#)

[IBM HPS POWER5 FLASH/Readme - Service Pack 3](#)

[IBM HPS POWER5 FLASH/Readme - Service Pack 2](#)

[IBM HPS POWER5 FLASH/Readme - Service Pack 1](#)

Health check/planning help

- [FLRT: Fix Level Recommendation Tool](#)
- [IBM Systems compatibility information](#)

- [Cluster support](#)

Changelog

ESSL/PESSL

History of problems fixed

Cumulative fixes thru ESSL 4.3.0.3

[August 12, 2008]

- PK67790 = Possible segmentation fault with large stride in FFT subroutines
- PK67796 = POWER5 performance improvements for DGER and DSPSV
- PK67963 = Improve matrix multiply performance for very large M
- PK67964 = POWER6 Performance Enhancements
- PK67965 = ESSL for AIX 430 PTF Set 3
- PK67972 = Possible segmentation fault for SGETMO
- PK68301 = Possible incorrect index from ICAMAX

Cumulative fixes thru ESSL 4.3.0.2

[December 18, 2007]

- PK55343 ESSL for AIX 430 PTF Set 2
- PK55345 Erratic DGEMM performance with 4KB pages on BladeCenter JS21
- PK55356 POWER6 Performance Improvements ENDSET

Cumulative fixes thru ESSL 4.3.0.1

[November 20, 2007]

- PK52576 Possible seg fault in SVES for $X=Y$ and $N > 3$ on AltiVec enabled processors
- PK52575 Possible wrong answer in CVEA when $X=Y=Z$ on AltiVec enabled processors
- PK52484 ESSL for AIX 430 PTF Set 1
- PK52483 Possible wrong answers from DGEADD and DGESUB when overlaying input and output

Cumulative fixes thru Parallel ESSL 3.3.0.2

[November 20, 2007]

- PK52541 Parallel ESSL for AIX 3.3 PTF Set 2
- PK52485 Possible incorrect results for PSRCFT2, PSRCFT2, PSCRFT2, PSCRFT3 on AltiVec enabled processors
- PK49859 Parallel ESSL header file not compatible with ESSL 4.3 header file
- PK49418 Support for MPICH-GM and C++ V8.0 on AltiVec enabled processors
- PK49414 Possible incorrect results for PSCFT3, PSRCFT3, PSCRFT3 on AltiVec enabled processors
- PK45573 Parallel ESSL for AIX 3.3 PTF Set 1

Cumulative fixes thru Parallel ESSL 3.2.0.1

[November 20, 2007]

- PK22151 pesslivp fails in 32-bit mode with MPICH-GM
- PK21664 Parallel ESSL for AIX 3.2 PTF Set 1
- PK21661 Parallel ESSL 3.2 BladeCenter JS21 and PPC970MP
- PK21656 Incorrect file permission for man page
- PK21648 Performance improvements in PDGEMV and PDSYMV
- PK21646 Null communicator issued by MPI_Abort

Cumulative fixes thru ESSL 4.2.0.14

[December 18, 2007]

- IZ06690 BUG IN LIBMPI THAT IS CAUSING THE HANG
- ppe.poe 4.2.2.14 U815155

Cumulative fixes thru ESSL 4.2.0.5

[November 20, 2007]

- PQ99011 ESSL IVP scripts need update for use with XLF 9.1
- PQ99009 POWER5 SMP performance improvements for p590/p595
- PQ99007 POWER5 performance improvements
- PQ99006 Performance improvements for DSPEV/SSPEV/ZHPEV/CHPEV/ DSYGV/SSYGV in SMP mode
- PQ99003 Possible wrong answers in DCFT3,SCFT3,DRCFT3,SRCFT3 in SMP mode with X=Y and non-matching increments
- PQ98997 Possible wrong answer in SPPSV and DPPSV after 2600 error message

- PQ98994 Incorrect error messages when using matrix inverse codes in SMP
- PQ98993 Performance improvements in DGEMM for small matrices
- PQ98991 Possible wrong answer in DSDCG
- PQ95435 ESSL V4.2 JS20 Support
- PK45568 ESSL 4.2.0 POWER6 Support
- PK45563 STRIDE returns incorrect values on POWER4, POWER5, PPC970
- PK25589 instfix returns incorrect output for ESSL 4.2.0 PTF Set 4
- PK25589 instfix returns incorrect output for ESSL 4.2.0 PTF
- PK20638 ESSL for AIX 420 PTF Set 5
- PK20634 Segmentation fault in packed L2 BLAS and packed LEQ subroutines
- PK19344 ESSL for AIX 420 PTF Set 4
- PK19343 ESSL 4.2 BladeCenter JS21 and PPC970MP support
- PK19342 Performance degradation for DCOSF and DSINF on POWER3 and POWER4
- PK19286 Improve SCFT3 performance for bad strides on POWER5
- PK18690 ESSL FFTs may seg fault for large arrays
- PK15144 Wrong answers with DGBMV on POWER3, M=N and $4 < ML > 50$
- PK14803 DGEMM performance on POWER4 when data not backed by large pages may be degraded
- PK14738 ESSL for AIX 420 PTF Set 3
- PK14737 Improve SIDE=L performance for ZTRMM, ZTRSM on POWER5
- PK14714 DGEMM performance improvements for POWER5+
- PK14713 POWER5 SMP performance improvements
- PK02186 ESSL for AIX 420 PTF Set 2
- PK02112 ESSL FFT subroutines may not provide bit identical answers in SMP
- PK02059 POWER5 performance improvements
- PK02057 POWER5 and JS20 performance improvements for DGEMM, SGEMM, CGEMM, ZGEMM

[Back to top](#)

Health check/planning help

- [FLRT: Fix Level Recommendation Tool](#)
- [IBM Systems compatibility information](#)

- [Cluster support](#)

Changelog

GPFS

Notes

Unless specifically noted otherwise, this history of problems fixed for GPFS applies to: fixes for GPFS 3.2.0 thru GPFS 3.2.0.1.

History of problems fixed

Cumulative fixes for GPFS 3.2.0 thru GPFS 3.2.0.1 (POWER5 HPS SP15 Addendum 1)

[November 20, 2007]

- IZ05802 CANT MOUNT IF ALL DISKS SUSPENDED
- IZ05802 LOG RECOVERY FAILS BECAUSE OF DELETED POOL STATUS
- IZ06367 AIX GPFS ALLOWED ADDING DISKS GREATER THAN 1TB
- IZ05802 POOLS WITH NO DISK NOT BEING CLEANED UP
- IZ05802 SLOW PERFORMANCE DURING TSM SCAN DUE TO MANY INDBLOCKS
- IZ05802 AIX HDISKS OVER 2TB NOT CORRECT SIZE
- IZ05802 Deadlocks possible if kxCommonReclock calls vnodeReleInternal
- IZ05802 CALLERS OF GETXATTR EXPECT -ERANGE (INSTEAD OF -ENOSPC)
- IZ05802 CLUSTER PAUSES WHILE ONE NODE SHUTS DOWN
- IZ05802 SAMBA KXGETSHARE CALLS EVENTUALLY DEADLOCK MACHINE
- IZ05802 Slow NFS write performance
- IZ05802 SIGSEGV IN ALLOCINODE WITH 2MB BLOCKSIZE AND NNODES 4096
- IZ05802 mmfsck error msg refers to previous msgs that do not exist
- IZ05802 nfs: inode marks dirty; pdflush stuck in sync_sb_inodes

Health check/planning help

- [FLRT: Fix Level Recommendation Tool](#)
- [IBM Systems compatibility information](#)

- [Cluster support](#)

Changelog

HPSNM

Notes

Unless specifically noted otherwise, this history of problems fixed for HPSNM applies to: HPSNM 1.4.1.x for AIX

History of problems fixed

Problems fixed in HPSNM 1.4.1.18

[August 24, 2007]

- IY96327 LSSWTOPOL fails to list frames when BPCS are
- IY96115 TOD cleanup code does not get reset when backing off
- IY94962 Rebuild TOD candidate list on Comm timeout/error
- IY94402 Error return from SWJTAGWRITE() stops that switches
- IY94336 FNM UTILITY: HPS_CHECKROUTE.PL is broken for 256-way
- IY94044 Cut AIX errlog entry for TOD thread exit
- IY93988 Reparents FNM GUI if CSM not installed
- IY92631 FIX FNM retry logic
- IY90547 Remove SmaGetLinkInfo() from early HW discovery
- IY90546 Close 2nd hole in SMALINKUP flow
- IY88538 HPS switch TOD enhancements
- IY87690 Timing hole in SwitchEvent credit count flow
- IY87688 FNM core dump with 16 SMA on p5 H
- IY87666 fnmd can start before hardware server is ready
- IY87606 cmd line lsswendpt et. al. fails on p5
- IY87506 Running HALT -Q when starting FNM - INITTAB
- IY86925 HPS switch board power cycle can cause packet loss
- IY86423 FNM_INIT fails to generate correct end point IDS
- IY86319 ELA bad location code
- IY85601 LSSWENDPT on the MS shows incorrect ROUTE_STATUS
- IY85104 Switch network hang on HPSNM restart after switch
- IY84139 p5 GA6 GUI end point view - verify link says SNM not

- IY84138 LogFile class check for bad bits when File System
- IY84137 Print out the fnm_init ack list from
- IY83210 Change credit count flow
- IY83205 JSB link enablement error
- IY81835 New adapter event not being formed by ELA
- IY81706 Core dump on injecting error in a 1-NSB cluster
- IY81450 hps_checkroute.pl does not handle large Endpoint IDs
- IY81448 create/add setup_service man page to install package
- IY81446 startFNMD will try to start fnmd even if not executable
- IY81445 Daemons need to change permission type on log files
- IY80470 HPS switch performance: FNM needs to validate link
- IY78362 TOD thread needs to read all available SMAS before
- IY76696 SNI location code is p4 format
- IY76165 Bad recovery from NSB isolate and rejoin
- IY76082 SNIs failed to configure on 16-way p5 IH
- IY75482 No error reported to SFP when ela backup is set
- IY74828 COMM layer should check it has read all the bytes in
- IY74065 SysEvent log entries slow down warm boot
- IY73981 Provide new message text to SFP that identifies when
- IY73970 Do not start tracelogd if csm.hpsnm is not installed
- IY73910 Missing serial numbers for the switches in LocCode
- IY73909 Improperly formed location codes for SNIs
- IY72988 LSSWTOPOL shows SVC REQ, but links are pingable
- IY72584 Wrap test for defective optical cable
- IY70554 Update to GUI documentation
- IY67240 Bad location codes passed to ELA

Health check/planning help

- [FLRT: Fix Level Recommendation Tool](#)
- [IBM Systems compatibility information](#)

- [Cluster support](#)

Changelog

HPS/SNI (Devices)

History of problems fixed

Unless specifically noted otherwise, this history of problems fixed for HPS/SNI applies to 1.2.0.x and 1.1.3.x for AIX.

Problems fixed in HPS/SNI 1.2.0.x (device support for AIX 53)

[August 12, 2008]

PTF 17

- IZ06443 SWITCH NTBL CHANGES UMASK FROM 022 TO 000
- devices.common.IBM.sni.ntbl 1.2.0.2 U820112

[June 19, 2008]

PTF 16

- IZ07145 HPS performance problem caused by depletion of the 64K PAGE POOL
- IZ23028 Add a 2 tuple hash option for TCP/IP traffic.
- IZ23031 Add sorting to the ML route table.
- devices.common.IBM.sni.rte 1.2.0.11
- devices.msg.en_US.common.IBM.sni.ml 1.2.0.1
- devices.chrp.IBM.HPS.rte 1.2.0.12
- devices.common.IBM.sni.ml 1.2.0.5

[June 19, 2008]

PTF 15

- IZ15886 MultiLink should check if underlying interface supports managed buffer
- IZ15886 HPS / IB sni.ml ptf15 co-existence

- devices.common.IBM.sni.ml 1.2.0.4

[August 24, 2007]

PTF 14

- IZ12720 infiniband support for ml
- devices.common.IBM.sni.ml 1.2.0.3
- IZ12720 minor changes to multi-link (ML) driver
- devices.common.IBM.sni.ml 1.2.0.3 U816069

PTF 13

- IZ09912 Changes for AIX61
- IZ06762 NTBL_RESOURCE_BUSY NTBL_CLEAN_WINDOW HPS_IOCTL DEADLOCK
- IZ09912 crash due to uninitialized variable
- devices.chrp.IBM.HPS.rte 1.2.0.11 U815638

PTF 12

- IZ04185 Fed. driver binds to a processor prohibiting SMT from disabling
- devices.chrp.IBM.HPS.rte 1.2.0.10 U813757

PTF 11

- IZ02259 cfigsni needs to check errno
- devices.common.IBM.sni.rte 1.2.0.9 U813015
- devices.chrp.IBM.HPS.rte 1.2.0.9 U813016

PTF 10

- IY95784 performance degraded when 64K large pages are used
- devices.common.IBM.sni.rte 1.2.0.8
- devices.chrp.IBM.HPS.rte 1.2.0.8

PTF 9

- IY93999 degraded performance mode after reboot
- IY93999 CHGSNI caused several nodes to crash and others non-responsive
- devices.chrp.IBM.HPS.rte 1.2.0.7 U811277

PTF 8

- IY92803 Remove idle forked server processes
- IY92801 PTF builds missing new VERIFICATION_TEST readme
- IY92801 Make the verification_test more fault tolerant
- devices.common.IBM.sni.rte 1.2.0.7 U810961

PTF 7

- IY91024 ERRPT MSG 284BDA76 incorrectly logged
- IY88872 Remove ifdef around open retry loop
- IY88077 P5 nodes crash when SNI_DD runs out of large pages
- devices.common.IBM.sni.rte 1.2.0.6 U810608
- devices.chrp.IBM.HPS.rte 1.2.0.6 U810609

[Back to top](#)

PTF 6

- IY86289 cfg error flow crash w/xmalloc_debug
- IY85351 RAS: SNI report ERRPT when insufficient large pages available
- devices.common.IBM.sni.rte 1.2.0.5 U809417
- devices.chrp.IBM.HPS.rte 1.2.0.5 U809418

PTF 5

- IY85178 GA8 update the hom_to_laddr map
- IY85177 RDMA lgpg BW ~6% lower on AIX 53E
- devices.common.IBM.sni.rte 1.2.0.4 U808009
- devices.msg.en_US.chrp.IBM.HPS.rte 1.2.0.2 U808010
- devices.chrp.IBM.HPS.rte 1.2.0.4 U808011

PTF 4

- IY83075 Export InfiniBand version of hal.h
- IY83075 64K page support
- IY83073 RCXT_READ_IOCTL
- IY82209 Preemption issues
- devices.common.IBM.sni.rte 1.2.0.3 U807736
- devices.chrp.IBM.HPS.rte 1.2.0.3 U807737

PTF 3

- IY80626 zdebugdrv fileset needs a post_u
- IY78180 SNIDD panic in XMEMDMA64()
- IY79629 CRASH IN MLTDD
- devices.chrp.IBM.HPS.rte 1.2.0.2 U806941
- devices.common.IBM.sni.ml 1.2.0.2 U806942

PTF 2

- IY78751 LL says rCxt are available even when not useable
- IY78751 CRASH IN SNI DEVICE DRIVER
- IY78653 Performance problem HPS W/ P575 SNI attached to DCM0 AND DCM1
- IY78649 SNI CONFIG failed during Cold Boot
- IY77427 Switch clog condition exposes IP FULL FIFO detection bug
- devices.common.IBM.sni.rte 1.2.0.2 U806323
- devices.chrp.IBM.HPS.rte 1.2.0.1 U806324

PTF 1

- IY77183 ntbl msg catalog says block instead of bulk transfer
- IY76048 increase time-out to address lapi_init() rc 608 failures
- devices.common.IBM.sni.rte 1.2.0.1 U805762
- devices.msg.en_US.common.IBM.sni.ntbl 1.2.0.1 U805763

[Back to top](#)

Problems fixed in HPS/SNI 1.1.3.x (Device support for AIX 5.2)

[June 19, 2008]

PTF Set 15

- IZ23059 HPS PERFORMANCE PROBLEM CAUSED BY DEPLETION OF THE 64K PAGE POOL
- devices.common.IBM.sni.rte 1.1.3.12
- devices.chrp.IBM.HPS.rte 1.1.3.14

PTF Set 14

- IZ02261 cfgsni needs to check errno
- devices.common.IBM.sni.rte 1.1.3.11 U813013

- devices.chrp.IBM.HPS.rte 1.1.3.13 U813014

PTF Set 13

- IY96858 Performance degraded when 64k large pages are used
- devices.common.IBM.sni.rte 1.1.3.10 U812543
- devices.chrp.IBM.HPS.rte 1.1.3.12 U812544

PTF Set 12

- IY94630 CHGSNI caused several nodes to crash. others non-responsive

PTF Set 11

- IY92804 Remove idle forked server processes
- IY92802 PTF builds missing new VERIFICATION_TEST readme
- devices.common.IBM.sni.rte 1.1.3.9 U810962

PTF Set 10

- IY91453 ERRPT MSG 284BDA76 INCORRECTLY LOGGED
- IY90971 P5 nodes crash when SNI_DD RUNS out of large pages
- IY89019 Remove ifdef around open retry loop
- devices.common.IBM.sni.rte 1.1.3.8 U810605
- devices.chrp.IBM.HPS.rte 1.1.3.10 U810606

PTF Set 9

- IY86290 cfg error flow crash w/xmalloc_debug
- IY86105 RDMA lgpg BW ~6% lower on AIX 53E
- IY86104 GA8 update the hom_to_laddr map
- IY85656 SNI report errpt when insufficient large pages
- devices.common.IBM.sni.rte 1.1.3.7 U809414
- devices.msg.en_US.chrp.IBM.HPS.rte 1.1.3.1 U809415
- devices.chrp.IBM.HPS.rte 1.1.3.9 U809416

PTF Set 8

- IY83077 64K page support
- IY83076 RCXT_READ_IOCTL
- IY82434 Preemption issues
- devices.common.IBM.sni.rte 1.1.3.6 U807738

- devices.chrp.IBM.HPS.rte 1.1.3.8 U807734

[Back to top](#)

PTF Set 7

- IY79998 CRASH IN MLTDD
- IY79550 SNIDD PANIC IN XMEMDMA64()
- devices.common.IBM.sni.ml 1.1.3.2 U806944
- devices.chrp.IBM.HPS.rte 1.1.3.6 U807130

PTF Set 6

- IY78648 SNI CONFIG FAILED during Cold Boot
- IY78744 increase time-out to address lapi_init()
- IY78750 LL says rCxt are available even when not usea
- IY77236 CRASH IN SNI DEVICE DRIVER
- IY78745 SWITCH CLOG CONDITION EXPOSES IP FULL
- IY78305 PERFORANCE PROBLEM HPS W/ P575 SNI ATTACHED
- devices.common.IBM.sni.rte 1.1.3.5 U806309
- devices.chrp.IBM.HPS.rte 1.1.3.5 U806310

PTF Set 5

- IY77182 ntbl msg catalog says block instead of bulk transfer
- devices.msg.en_US.common.IBM.sni.ntbl 1.1.3.1 U805764

PTF Set 4

- IY74189 HPS ON p5 DD CALL HANG
- IY73996 need a tool to easily tell how much
- devices.common.IBM.sni.rte 1.1.3.4 U804896
- devices.chrp.IBM.HPS.rte 1.1.3.4 U804897

PTF Set 3

- IY72510 add alias4 support for ML interface
- IY69913 ALIAS support for Federation
- devices.common.IBM.sni.rte 1.1.3.3 U804349
- devices.msg.en_US.common.IBM.sni.ml 1.1.3.1 U804350
- devices.chrp.IBM.HPS.rte 1.1.3.3 U804351

- devices.common.IBM.sni.ml 1.1.3.1 U804352

PTF Set 2

- IY71293 print SRC's in parser log
- devices.common.IBM.sni.ntbl 1.1.3.1 U803736
- devices.common.IBM.sni.rte 1.1.3.2 U803715
- devices.chrp.IBM.HPS.rte 1.1.3.2 U803716

PTF Set 1 (Plus cumulatives fixes for HPS Devices 1.1.x thru 1.1.3.1)

- IY70551 checkstop during mpfatal recovery scomm collision
- IY70294 Required fixes for devices 113
- IY65208 table and slot allocation
- IY64728 need to close the dump window
- IY64726 p5 mp_fatal dump support part 1
- IY64724 handle H_BUSY from dealloc_iotranslate
- IY63326 Fixes for HPS/SNI
- IY63023 HAL snap can be dumping user data
- IY63022 return canp_fifo_rdma as css_dev type on POWER 5
- IY62838 No adapters 'Available' after rebooting LPAR
- IY62834 dd needs to kill all bad window procs,
- IY62811 Loopback driver and tools
- IY62809 fix exposure in system();
- IY62808 trace the lmt for window fatals in close window
- IY62806 TCEs should be traced at error level
- IY62803 RDMA: Workaround for Austin AIX defect #470058.
- IY62769 Implement the re-cable flows
- IY62765 Ship additional debug tools.
- IY62764 updates for new window fatal condition (rCxt ID)
- IY62762 RDMA: Enable odm entry and chgsni support for
- IY62760 investigate tcb support.
- IY62759 fix 109821
- IY62755 re-work getsmadata for squadrons - new prd signature
- IY62754 Network ID support for mixed p4/p5
- IY60102 Change default network mapping for 2 plane sy
- IY60015 Alias Support for Federation
- IY58914 Add dynamic pool reduction in IP support
- IY58912 Test case not getting window fatal notification
- IY58559 Error data logging for window fatals.
- IY58556 swclock for can_fifo_rdma css_devtype
- IY58525 FED RDMA panic in ml_output/sn_xmit_spkt
- IY58524 core dump in hps_snap in handling MP_DOWN

- IY58523 Set and reset read_pending in dgsp_read
- IY58522 dump_key value reset
- IY58521 H_CALL perf timers
- IY58124 Required Fixes for devices 1.1.1.0
- IY55499 wrong command in sni.snap for ML
- IY55381 PDAT for state set to default down
- IY55282 iptrace crashes node with HPS switch
- IY53788 Workaround for hardware defect
- IY52058 DD code to reset the FTM6 send sequence number table
- IY51425 Undo Path Table recovery change
- IY49654 FED ucode loader interface change
- devices.common.IBM.sni.ntbl 1.1.3.0 U803366
- devices.msg.en_US.chrp.IBM.HPS.rte 1.1.3.0 U803367
- devices.common.IBM.sni.ml 1.1.3.0 U803368
- devices.common.IBM.sni.rte 1.1.3.1 U803544
- devices.msg.en_US.common.IBM.sni.ml 1.1.3.0 U803371
- devices.msg.en_US.chrp.IBM.HPS.hpsfu 1.1.3.0 U803372
- devices.msg.en_US.common.IBM.sni.rte 1.1.3.0 U803373
- devices.chrp.IBM.HPS.rte 1.1.3.1 U803545
- devices.chrp.IBM.HPS.hpsfu 1.1.3.0 U803375
- devices.msg.en_US.common.IBM.sni.ntbl 1.1.3.0 U803376

[Back to top](#)

Health check/planning help

- [FLRT: Fix Level Recommendation Tool](#)
- [IBM Systems compatibility information](#)

- [Cluster support](#)

Changelog

Parallel Environment

Notes

Unless specifically noted otherwise, this history of problems fixed for PE applies to PE 4.2.x and PE 4.3.x.

History of problems fixed

Cumulative fixes thru PE 4.3.2.3

[August 12, 2008]

- **IZ24833 PREEMPTION FAILURE - POE FAILS WITH 0031-463**
- **IZ24517 POE PRIORITY fails due to UID check failing**
- **IZ24192 LAUNCH OF 6400 MPI TASK JOB FAILS**
- **ppe.poe 4.3.2.3 U819647**

Cumulative fixes thru PE 4.3.2.2

[August 12, 2008]

- **IZ21566 update cupset_query man page**
- **IZ21566 problem with MP_TASK_AFFINITY=MCM on SLES10**
- **ppe.poe 4.3.2.2 U818935**
- **ppe.man 3.3.2.1 U818934**

Cumulative fixes thru PE 4.3.2.1

[June 19, 2008]

- **IZ13077 MPI_SCATTER SCALING PERFORMANCE**
- **IZ16253 MP_STATISTICS explanation incorrect**
- **IZ16253 MP_PMDLOG_DIR/-pmdlog_dir Does not write full log**

- **IZ16253 malloc failure when running at scale**
- **IZ16253 Adjust QP data structure allocations**
- **IZ16253 PE Install Fail on X Sles 10**
- **IZ16253 POE incorrectly detects LoadL 3.4.2 or earlier**
- **IZ16253 MP_PRINTENV inverts MP_SINGLE_THREAD value**
- **IZ16253 update VRMF of ppe_rcha and rsct_rni**
- **IZ16253 NOPOLL wait in MPI uses CPU cycles**
- **ppe.poe 4.3.2.1 U818010**

Cumulative fixes thru PE 4.3.1.5

[November 20, 2007]

- **IZ04217 change vrmf for PE and LAPI.**
- **IZ04217 PE 4.3.1 AIX GA package is broken.**
- **IZ04217 MPI: Core dump with procs>5 when calling MPI_Wait/Test.**
- **IZ04217 Extend MP_TASK_AFFINITY to CORE and CPU.**
- **IZ04217 AIX LAPI244 & PE431 post install broken.**
- **IZ04217 MPI: Program hangs when calling MPI_Win_call_errhandler.**
- **IZ02784 MPI IO PROBLEM WITH COMMUNICATOR OTHER THAN MPI_COMM_WORLD.**
- **IZ02781 POE SAMPLE NRT USAGE ENHANCEMENTS.**
- **IZ01193 FALSE POE LARGE PAGE WARNING WHEN -BDATAPSIZE OPTION IS USED.**

Cumulative fixes thru PE 4.3.1.3

- **IZ01145 ADD TIMESTAMPS FOR STARTUP DEBUGGING**
- **IY99870 MPI Internal error when specifying File I/O.**
- **IY99869 Ckpt jobs hang in mpi_ckpt_thread_handler with PPE.**
- **IY98217 Incorrect msg matching for short MPI_Irecv from.**
- **IY98216 MPI_TYPE_CREATE_RESIZED IN 64 BITS CAN SEGFAULT.**
- **IY98133 PMB BENCHMARK PROBLEM WITH MPI.**
- **IY97393 SMT processor core affinity support phase 1.**
- **IY91238 Excessive LAPI Timer pops.**
- **IY89226 RDMA lgpg BW ~6% lower on AIX 53E.**
- **IY85131 mpi/us and lapi/ip job fails lapi rc=265155268.**
- **IY85109 SW MP_DOWN reason 2d on AIX 5.2M doing checkpoint.**
- **IY83589 LAPI must handle GS_REOPEN_FAIL.**
- **IY82966 MPI_Init() seg fault w/ multiple instances on.**
- **IY81915 RCXT_READ_IOCTL.**
- **IY78479 LAPI interrupts disabled when MP_MSG_API=MPI,LAPI.**

- **IY77180 Termination hang with 64 bit trace library.**
- **IY74062 Striping rDMA performance.**
- **IY72866 LAPI_Init/Term timing problem.**
- **IY71910 RDMA_SETUP FAILURE IN LAPI.**
- **IY71295 Remove the use of LAPI kernel extension timer.**
- **IY70553 Required fixes for LAPI 233.**

Cumulative fixes thru PE 4.2.2.15

- **IZ13560 WHEN MP_PRIORITY IS USED THEN CTRL-D IS NOT PASSED IN STDIN**
- **ppe.poe 4.2.2.15 U817194**

Cumulative fixes thru PE 4.2.2.14

- **IZ06690 BUG IN LIBMPI THAT IS CAUSING THE HANG**
- **ppe.poe 4.2.2.14 U815155**

Cumulative fixes thru PE 4.2.2.13

- **IZ00125 FALSE POE LARGE PAGE WARNING WHEN -BDATAPSIZE OPTION IS USED**
- **IZ02644 MPI IO PROBLEM WITH COMMUNICATOR OTHER THAN MPI_COMM_WORLD**
- **ppe.poe 4.2.2.13 U814149**

[Back to top](#)

Cumulative fixes thru PE 4.2.2.12

- **IZ00901 ADD TIMESTAMPS FOR STARTUP DEBUGGING**
- **ppe.poe 4.2.2.12 U812727**

Cumulative fixes thru PE 4.2.2.11

- **IY95969 MPI_TYPE_CREATE_RESIZED IN 64 BITS CAN SEGFAULT**
- **IY97590 Incorrect msg matching for short MPI_Irecv from MPI_ANY_SOURCE**
- **ppe.poe 4.2.2.11 U811973**

Cumulative fixes thru PE 4.2.2.10

- **IY96112 CHECKPOINT/RESTART GIVES WRONG RESULT INTERMITTANTLY**

- **ppe.poe 4.2.2.10 U811536**

Cumulative fixes thru PE 4.2.2.9

- **IY86318 MPXLF_R USES THE HEAD COMMAND THAT IS NOT FU**
- **IY94251 Fix for -tlp requires change**
- **ppe.poe 4.2.2.9 U811264**

Cumulative fixes thru PE 4.2.2.8

- **IY93106 PMD does not recognize LoadL 3.4 for affinity**
- **IY93258 Occasional race condition error in value returned by MPI_Iprobe**
- **ppe.poe 4.2.2.8 U811018**

Cumulative fixes thru PE 4.2.2.7

- **IY87217 PCT HANGS AT THE END OF THE EXECUTION**
- **IY92021 Retrofit fix for large page warning/kill capability**
- **IY89417 PREMATURE PULSE TIMEOUTS WITH REDIRECTED STDIN AND CHECKPOINT**
- **IY89890 BUG IN LAPI_DGSM_GOSUB COMMAND OF DGSP GENERATED BY MPI_GATHER**
- **IY92022 Miscellaneous NTBL Sample flag cleanup**
- **IY91676 File /tmp/s.pedb.<poe_pid>.<uid> has permission 777**
- **IY92071 MPCIE LOCK LEAK PROBLEM**
- **ppe.perf 4.2.2.3 U810901**
- **ppe.poe 4.2.2.7 U810902**

Cumulative fixes thru PE 4.2.2.6

- **IY87218 DOCUMENTATION NEEDED ON DCPLS AFFECT ON 64-BIT FORTRAN CODE**
- **IY87778 INCORRECT MESSAGE SIZE RETURNED FROM MPI_PROBE**
- **IY88058 MEMORY LEAK IN MPI COMPONENT**
- **IY88225 THE STORAGE USED FOR DERIVED DATA TYPE IS NOT FREED**
- **ppe.perf 4.2.2.2 U809969**
- **ppe.poe 4.2.2.6 U809970**

Cumulative fixes thru PE 4.2.2.5

- **IY83411 HANDLE MORE EFFICIENTLY REDIRECTED STDIN**
- **IY83782 TOTALVIEW BROKEN WITH PE 4.2.2**

- **ppe.poe 4.2.2.5 U808958**

Cumulative fixes thru PE 4.2.2.4

- **IY84543 Linux: Memory Leak in MPI_Datatype create/free cycle**
- **IY82525 MPI_BCAST() FAILS WITH COMMUNICATION TIMEOUT**
- **IY82525 RCXT_READ_IOCTL**
- **IY83126 MPI_TEST SOME AND MPI_WAIT SOME FAIL.**
- **IY84171 MPI_CART_CREATE FAILS**
- **ppe.poe 4.2.2.4 U807926**

[Back to top](#)

Cumulative fixes thru PE 4.2.2.3

- **IY80928 llq -l and llsummary -l does not display rs**
- **IY80928 llq -l does not show adapters as not shared**
- **IY79048 MP_PRINTENV PRINTS MISLEADING INFO ABOUT MP_INSTANCES**
- **IY79461 FAILURE TO CREATE LIGHT WEIGHT CORE FILES.**
- **IY80929 Update AIX Trace hooks in MPC I**
- **IY81148 INCORPORATE NEW LAPI ENV VARIABLES IN NEXT**
- **IY81379 timing issues in rset attachment**
- **ppe.man 4.2.2.2 U806968**
- **ppe.poe 4.2.2.3 U806969**

Cumulative fixes thru PE 4.2.2.2

- **IY78905 SAMPLES SCRIPT /USR/LPP/PPE.POE/SAMPLES/AFS/**
- **IY79190 POE GENERATES NULL CHARACTERS ON MP_PRIORITY**
- **IY79312 POE FAILS WHEN TOTALVIEW FAST CONDITIONAL BR**
- **IY79423 THE PARTITION MANAGER DAEMON (PMDV4) TERMINA**
- **IY80314 MP_BUFFER_MEM second argument being ignored**
- **ppe.poe 4.2.2.2 U806683**

Cumulative fixes thru PE 4.2.2.1

- **IY78110 Problems with probe insertion when execute**
- **IY78397 ASCII SWL: POE's _us_info() does not scale**
- **IY75356 MP_WAIT_MODE=nopoll not working as designed**
- **IY78110 Linux: POE Segmentation fault with option -sa**
- **IY78398 CHANGE 0031-833 MESSAGE SETTING FROM ATTN TO**
- **IY78110 bg_node_mode=VN: llq shows Mode VN llstatus -**

- **IY76953 MP_STDINMODE 1 FAILS WHEN 2 TASKS**
- **IY78110 BG: 3 problems for BlueGene specific API**
- **IY78110 Linux: MPI_File_open coredump when MP_EUIDEVE**
- **IY78110 Linux: EMPTY_CTLIN_PIPE macro causes MPI_Abor**
- **IY78110 Missing RDMA error messages in catalog**
- **IY78110 xmem attach failure on single element vector**
- **IY78110 Probetype hwcount missing for Power5 node**
- **IY78260 DELAY TIME WHEN POE MUST RETRY A SOCKET CONN**
- **IY78529 CLOCK DRIFT WITH POE CO-SCHEDULER**
- **ppe.perf 4.2.2.1 U806315**
- **ppe.man 4.2.2.1 U806316**
- **ppe.poe 4.2.2.1 U806317**

4.2.2.0 (Mod level refresh PTFs)

- **ppe.perf 4.2.2.0 U805892**
- **ppe.pvt 4.2.2.0 U805893**
- **ppe.man 4.2.2.0 U805894**
- **ppe.poe 4.2.2.0 U805895**

[Back to top](#)

Health check/planning help

- [**FLRT: Fix Level Recommendation Tool**](#)
- [**IBM Systems compatibility information**](#)

- [Cluster support](#)

Changelog

VSD

Notes

Unless specifically noted otherwise, this history of problems fixed for VSD:

Cumulative VSD 4.1.0.x fix list for AIX52 and AIX53.

History of problems fixed

PTF 28

- IZ19828 prReadKeys not found in vsd.reserve
- rsct.vsd.rvsd 4.1.0.19

PTF 27

- IZ16088 RVSD recovery hanging
- IZ17043 system panic
- rsct.vsd.vsdd 4.1.0.23
- rsct.vsd.rvsd 4.1.0.18

PTF 26

- IZ05817 RVSD code tightening
- IZ05816 VSD read I/O is blocked on client node
- rsct.vsd.vsdd 4.1.0.22 U815454
- rsct.vsd.rvsd 4.1.0.17 U815455

PTF 25

- IY99369 vsds may not activate
- IY98321 IOCINFO ioctl not correct for large vsds
- rsct.vsd.vsdd 4.1.0.21 U812842

- rsct.vsd.rvsd 4.1.0.16 U812843

PTF 24

- IY94854 vsd.snap needs additional info
- IY94127 System Crash in CpyRd()
- rsct.vsd.vsdd 4.1.0.20 U811224
- rsct.vsd.cmds 4.1.0.18 U811225

PTF 23

- IY92553 defvsd prints out extraneous info
- IY92552 removed -fa failed to cleanup vgs on VSD primary
- IY92200 Enhance debug/data integrity code
- IY89863 RVSD hangs during low paging space
- rsct.vsd.vsdd 4.1.0.19 U810964
- rsct.vsd.rvsd 4.1.0.15 U810965
- rsct.vsd.cmds 4.1.0.17 U810966

PTF 22

- IY88936 System crash in RcvRmt and RcvCtl
- IY87533 lower RcvSigRdyFlg retry time
- rsct.vsd.vsdd 4.1.0.18 U809977

PTF 21

- IY86218 vsd hangup
- IY86217 Device support enhancements
- IY86216 Invalid recv message handling
- rsct.vsd.vsdd 4.1.0.17 U808926
- rsct.vsd.cmds 4.1.0.16 U808927
- rsct.vsd.rvsd 4.1.0.14 U808928

PTF 20

- IY85001 VSD data corruption on READS
- rsct.vsd.vsdd 4.1.0.16 U808007

PTF 19

- IY83647 panic:SndRmt vsdkt threadentry

- IY83353 VSD system panic in RcvRmt
- rsct.vsd.vsdd 4.1.0.15 U807742

PTF 18

- IY79818 Trace gathering fix
- IY78902 VSD 2-node recovery is failing
- IY78343 Node panic in vsd during failover in VSDD:SUSVSD
- IY78175 VSD.SNAP error in the subroutine GET_CSS_STUFF
- IY78174 VSD.SNAP gets error when remote FS NAMED /TMP IS
- rsct.vsd.vsdd 4.1.0.14 U806617
- rsct.vsd.vsdrm 4.1.0.7 U806618
- rsct.vsd.cmds 4.1.0.15 U806619
- rsct.vsd.rvsd 4.1.0.13 U806620

PTF 17

- IY76888 vsd server crash
- IY76538 add server rdrsp throttle
- rsct.vsd.vsdd 4.1.0.13 U805754
- rsct.vsd.vsdrm 4.1.0.6 U805755
- rsct.vsd.cmds 4.1.0.14 U805756
- rsct.vsd.rvsd 4.1.0.12 U805757

PTF 16

- IY73538 VSD read/write throttle
- rsct.vsd.vsdd 4.1.0.12 U804742
- rsct.vsd.vsdrm 4.1.0.5 U804743
- rsct.vsd.cmds 4.1.0.13 U804744
- rsct.vsd.rvsd 4.1.0.11 U804745

PTF 14

- IY71399 I/O hung on cvsd disks after pulled all fc cables
- rsct.vsd.vsdd 4.1.0.11 U803737

PTF 13

- IY70282 stoprpdomain from an online node fails to stop
- IY66940 Disallow IP fragmentation
- IY66940 missed an error path

- rsct.vsd.vsdd 4.1.0.10 U803536
- rsct.vsd.vsdrm 4.1.0.4 U803537
- rsct.vsd.cmds 4.1.0.12 U803538
- rsct.vsd.rvsd 4.1.0.10 U803539

PTF 11

- IY64262 PR capable disks won't vary online
- IY64260 VSD needs to reset Persistent Reserve attributes
- IY64259 Issues with file permissions
- rsct.vsd.cmds 4.1.0.11 U801290
- rsct.vsd.rvsd 4.1.0.9 U801291

PTF 10

- IY63020 Accidental rpd removal destroys vsd config
- IY62693 cfgvsd should fail if bad IP addresses
- IY62692 with domain down vsdata1st -n returns no error
- IY62691 IBM.vsdd and rvsdd must handle SIGDANGER
- IY62690 VSD Multi-cluster Support
- IY62680 ctlvsd -t displays header out of order
- rsct.vsd.vsdd 4.1.0.9 U800691
- rsct.vsd.vsdrm 4.1.0.3 U800692
- rsct.vsd.cmds 4.1.0.10 U800693
- rsct.vsd.rvsd 4.1.0.8 U800694

PTF 9

- IY60095 Need VSD support for greater than 2TB LUNs
- IY59982 vsd error messages not getting displayed prop
- rsct.vsd.vsdd 4.1.0.8 U800292
- rsct.vsd.cmds 4.1.0.9 U800293

Health check/planning help

- [FLRT: Fix Level Recommendation Tool](#)
- [IBM Systems compatibility information](#)

- [Cluster support](#)

Changelog

LAPI

Notes

Unless specifically noted otherwise, this history of problems fixed for LAPI applies to LAPI 2.3.3 thru LAPI 2.4.5.x:

History of problems fixed

Cumulative fixes for LAPI 2.4.4 thru LAPI 2.4.5.3

[August 12, 2008]

- IZ24518 Create an nrt_clean
- IZ24518 Segv in hal_ib_open during task scaling
- IZ24518 Hang in cleanup after errs from lapi_rc_rdma_verbs_wrappers
- IZ24518 PNSD discards batch updates if size greater than 2048

Cumulative fixes for LAPI 2.4.4 thru LAPI 2.4.5.2

[August 12, 2008]

- IZ21569 Failover/Recovery not working on JS22 Linux/IB.
- IZ21569 stripe hal_open Seg Fault on 4 network IB cluster for NIC.

Cumulative fixes for LAPI 2.4.4 thru LAPI 2.4.5.1

[June 29, 2008]

- IZ18193 seg fault resuming bulkxfer mpi,lapi job
- IZ18193 LAPI asserts seen running Parallel ESSL bucket
- IZ18193 intermittent seg fault preempting IP job
- IZ18193 Fix for PNSD affinity per job info
- IZ18193 AIX 5.3N hal.ib cores in free_y on JS21 Cisco IB

- IZ18193 Introduce LAPI_DEBUG_4KMTU_RDMA environment variable
- IZ18193 Performance Problem Fix with Shared Memory

[January 24, 2007]

- IZ03658 LAPI 2.4.4 MISSING /USR/LIB links for new files on PTF upgrade.

Cumulative fixes for LAPI 2.4.4 thru LAPI 2.4.4.3

- IZ04710 LAPI 2.4.4 AIX GA package is broken.
- IZ04710 AIX LAPI244 & PE431 post install broken.

LAPI 2.4.4 thru LAPI 2.4.4.2

[November 15, 2007]

- IY99871 lapi_init RC605 caused by pnsd exit due to SIG PIPE.
- IY99871 add new env HAL_ENABLE_ERR_PRINT to print more info for rc=605.
- IY99871 Test lapi_vec_gen_64_2 segfaults on LAPI_Term.
- IY99871 TI build breaks during building stage on AIX.
- IY99871 Seg fault in fmemcpy.icopy.
- IY99871 MP_UDP_PACKET_SIZE should default to 2044
- IY99871 MPI not finalized errors while running NIC bucket
- IY99871 Linux: LAPI TI PTF2 RPM VRMF is off by 1.
- IY99871 IB PTF2: reduce receive FIFO size to 4 MB.
- IY99871 Add Federation style completion handling for performance.
- IY97394 lapi_init RC605 and seg fault on sles IB.
- IY97394 Tune P6 copy routine.
- IY97394 Linux: PE 431 & LAPI 244 Fix Level update.
- IY96860 Aix Partition went to KDB.
- IY96358 Updates for LAPI 2.4.4.0.
- IY96003 dgsm does not handle 0 reps LAPI_DGSM_GOSUB command.
- IY94000 CHANGE LAPI_ERR_INTERNAL5(876) LAPI_ERR_REG_TIMER(414).
- IY92594 LAPI LOCK PROBLEM DURING PREEMPTION.
- IY90799 INCREASE DEFAULT TIMEOUT FOR LAPI INIT.

LAPI 2.3.3 thru LAPI 2.3.3.10

[June 21, 2008]

- IZ13743 LAPI Samples don't print error message
- IZ15308 MPI IO MPI_WRITE_AT_ALL ON 1 TASK USING IP HAS POOR PERFORMANCE

[November 15, 2007]

- IZ05684 Seg fault in fmemcopy.icopy

Cumulative fixes for LAPI 2.3.3 thru LAPI 2.3.3.8

[November 15, 2007]

- IY91238 Excessive LAPI Timer pops.
- IY89226 RDMA lpgg BW ~6% lower on AIX 53E.
- IY85131 mpi/us and lapi/ip job fails lapi rc=265155268.
- IY85109 SW MP_DOWN reason 2d on AIX 5.2M doing checkpoint MPI RDMA job.
- IY83589 LAPI must handle GS_REOPEN_FAIL.
- IY82966 MPI_Init() seg fault w/ multiple instances on IP.
- IY81915 RCXT_READ_IOCTL.
- IY78479 LAPI interrupts disabled when MP_MSG_API=MPI,LAPI.
- IY77180 Termination hang with 64 bit trace library.
- IY74062 Striping rDMA performance.
- IY72866 LAPI_Init/Term timing problem.
- IY71910 RDMA_SETUP FAILURE IN LAPI.
- IY71295 Remove the use of LAPI kernel extension timer for SHM.
- IY70553 Required fixes for LAPI 233.

Health check/planning help

- [FLRT: Fix Level Recommendation Tool](#)
- [IBM Systems compatibility information](#)