

IBM HPS POWER5 Readme - Service Pack 12

[Introduction](#)

[Supporting documentation](#)

[Component updates](#)

[Installation sequence](#)

[AIX LP](#)

[Known problems and workarounds](#)

[Fix list](#)

[Readme archive](#)

Updated 2/19/2007

Introduction

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) LP
- Cluster Systems Management (CSM) LP
- Reliable Scalable Cluster Technology (RSCT) software
- AIX 5L Version 5.2 and 5.3
- Various supporting AIX LPs

This Service Pack provides support for release SF235_214 / BP235_212.

Important note:

Power Subsystem update BP235_212 should only be installed on p5 systems using the IBM pSeries High Performance Switch and is downloadable from a the special links provided in the [Component Update](#) section of this Readme. Obtain the System Firmware update SF235_214 from the normal IBM Microcode Download site location.

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, you must follow the workaround as described in the updated HMC [Installation Notes](#) section of this ReadMe.

HPS p5 SP11 allows use of the HPS TOD "Time of Day" function (refer to "Restrictions" section).

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.

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 (part of HPSNM/FNM)

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 (part of HPSNM/FNM)
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 + MH00839
IBMhsc.HPSNM_log-1.4.1.12-1 RPM (part of HPSNM/FNM)
IBMhsc.NM_Common-1.0.2.0-1

CEC: Global Firmware (GFW) - Version SF 240_284
FRAME: Power Code (BPC code) - Version BP240_203

CSM MS: AIX 5L Version 5.2 TL 09
AIX 5L Version 5.3 TL 05
HPSNM/FNM - CSM.HPSNM 1.4.1.16
RSCT -2.3.10 (AIX52) / 2.4.6 (AIX53)
CSM - 1.5.1

LPARS: [HPS/CSS PTFs \(LPARs\)](#) -

AIX 5L Version 5.2/ CSS 1.1.3
AIX 5L Version 5.3/ CSS 1.2.0

Supported LP's: as listed in [Detailed LP level check](#).

SP5 introduced support for 240 firmware.

New functions and features

The following items are new functions and features:

- Support for the model 9116-561.
- Support for OpenPower p5-511 and p5-511Q.
- Support for pSystem 521, pSystem 521Q, and pSystem 551.
- 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 the "quiet office" acoustic insulation package on pSystem 521 and pSystem 521Q servers.
- Support for a thin console on iSeries systems.
- 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 is increased to six on an iSeries model 595.

SP3 introduced a modified install plan intended to reduce the maintenance window down time.

Overview

Service Pack upgrade procedure to reduce the maintenance window for the service packs upgrades.

Procedure:

1. 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.
2. On the HMC - install / upgrade:
 - Upgrade ELA master and wait until ELA master is up and running.
 - Upgrade the rest of the HMCs.
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.
4. GFW upgrade:
 - Disable SNM.
 - If the new GFW supports concurrent upgrade; then select concurrent upgrade during the update process.
 - If the new GFW don't supports 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 don't have to activate Power Code or GFW; then reboot the LPARs to activate alt disk install.
 - If you don't 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.

Note:

This SP updates both the GFW and the BPC code concurrently.

If you are coming from SP2 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 Update/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 Update/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.

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.

IBM HPS POWER5 Readme - Service Pack 12

[Introduction](#)

[Supporting documentation](#)

[Component updates](#)

[Installation sequence](#)

[AIX LP](#)

[Known problems and workarounds](#)

[Fix list](#)

[Readme archive](#)

Supporting Documentation

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](#)
- [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](#)

IBM HPS POWER5 Readme - Service Pack 12

[Introduction](#)[Supporting documentation](#)[Component updates](#)[Installation sequence](#)[AIX LP](#)[Known problems and workarounds](#)[Fix list](#)[Readme archive](#)

Component Update/Download Information

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

Service pack update and download information

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. See Notice regarding 7310-C04 and 7310-C05/ internal modem compatibility issue.	HMC Support Select "HMC 4.5" Downloads for Updates are on the "Corrective service" tab. Click the "View" link for an HMC Update package to see the Readme.
	235 HMC V5R1 Select "HMC_Recovery_CD_V5R1.0" This release requires that the HMC is installed with HMC V5R1.0. PTF MH00593 is an HMC Required Maintenance Package (HMC Version 5, Release 1.2.). This Maintenance Fix <i>must</i> be installed on HMC Version 5 Release 1.0. 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.	HMC Support Select "HMC 5.1" 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.

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.

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

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

Click the "View" link for an HMC package to see the Readme.

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

[HMC corrective service](#)

240

HMC V6R1.0 + MH00839

Select "HMC 6.1"

This release requires that the HMC is installed with HMC V6R1.0.

Under "HMC corrective service support," select "HMC V6.x" under " Select HMC 6.1"

Note: MH00839 is required to get to V6R1.1

Note: You must use Recovery media or the Network to perform an Upgrade to update to HMC Version 6 Release 1.0.

For the HMC ReadMe Scroll to "Update corrective service"

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

Select "view" In the "ReadMe" column

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

240: Fix pack BP240_203
Fix pack SF240_284

**Power Subsystem
Microcode**

[Microcode Updates](#)

AND

System Firmware

****Note:**

Power Subsystem update BP235_212 should only be installed on POWER5 systems using the IBM pSeries High Performance Switch. Download this update by clicking on the following links:

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

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

5200-09-02

bos.mp64: 5.2.0.97

[pSeries support fixes](#)

AIX53

5300-05-03

bos.mp64: 5.3.0.53

[AIX V 5.2](#)

[AIX V 5.3](#)

AIX

Also Apply Critical fixes as appropriate.

[AIX 5L Critical fixes](#)

Note:

Critical fixes for AIX 5L service are no longer being updated.

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 5L Service Strategy and Best Practices](#).

To download Service Packs for AIX visit [Quick links for AIX fixes](#).

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**
 - conserver-8.1**

[openCIMOM](#)

openCIMOM
(CSM requisite)

openssl-0.9.7d-2

[sourceforge.net -
openssh-aix](#)

SSL_
(CSM requisite)

AIX-rpm-5.2.0.40-1

[sourceforge.net -
openssh-aix](#)

SSH_
(CSM requisite)

openssh-3.8.1p1_52.tar.gz

On the CSM MS:

This release requires HPSNM 1.4.1.16
apply HPSNM AIX PTF U810662, APAR IY88538

[pSeries support fixes](#)

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.

[Fix Central](#)

HPSNM

On the HMC:

230 Level:
HMC v 4.5
HPSNM Service Pack 3 MB01192 MH00401

235 Level
HMC v 5.1
HPSNM Service Pack 4 MB01237 MH00461

240 Level:
HMC v 6.1
HPSNM Service Pack 7 MB01648 MH00817

[Hardware](#)
[Management Console](#)

Support for UNIX
servers and Midrange
servers

Select an HMC
release, and then select
the "HPSNM/IBMNM
fixes" tab

IBM HPS POWER5 Readme - Service Pack 12

[Introduction](#)[Supporting documentation](#)[Component updates](#)[Installation sequence](#)[AIX LP](#)[Known problems and workarounds](#)[Fix list](#)[Readme archive](#)

Installation Sequence

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
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 these IBM Documents:

- [pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide for IBM eServer p5 servers](#)
- [CSM for AIX5L V1.4.1 Planning and Installation Guide](#)
- [CSM for AIX5L V1.4.1 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.

Install HMC

Install Hardware Management Console (HMC) software on all HMCs in the cluster.
There are no required efixes to apply.

Documentation Reference:

[Hardware Management Console](#)[Hardware Management Console V4R5.0 Readme](#)[Hardware Management Console V5R1.0 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 Power5servers 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, ensure that the HMC resources in the ClusterPeerDomain that were created with addpeer have the Manager_Configured attribute set to 1 prior to upgrading.

On the CSM management server enter "lsrhws -m". This will list all of the HMCs that are defined in the ClusterPeerDomain. The output for one HMC would be similar to:

```
lsrhws -m
-----
Manager_Type           = "HMC"
Manager_IP_A           = "20.0.0.5"
Manager_IP_B           = ""
Manager_Name           = "40.0.0.5"
Manager_MTMS           = "7315CR2*10407DA"
Manager_Configured     = 0
```

If the HMC listed above was added to the ClusterPeerDomain using addpeer and the "Manager_Configured" attribute equals "0" then it needs to be changed to "1".

If all of the HMCs defined in the ClusterPeerDomain were added with addpeer, then all HMCs may be updated at once using the following command:

```
chrhws -s 'Manager_Type=="HMC"' -a Manager_Configured=1
```

If only one HMC resource needs to be updated, then use the following command:

```
chrhws -s 'Manager_Name=="10.0.0.71"' -a Manager_Configured=1
```

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: 1.1
HMC Build level 20061103.1
```

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: 1.1

HMC Build level 20061103.1

Required fixes for V6R1.1

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

[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.

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 Service Guide 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

[For levels refer to the Component Update/Download Information for 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.

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.

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](#).

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-1.ibm.com/servers/aix/products/aixos/linux/download.html>
- The following opensource images also available on the AIX Toolbox page, <http://www-1.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-03.ibm.com/servers/eserver/support/unixservers/aixfixes.html>

Notes:

- i. 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.
 - ii. Check that the size of the /var directory on the CSM Management Server is approximately 1 Gigabyte to accommodate the HPSNM log files.
 - iii. 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:

- o lspp -L csm.client csm.core csm.diagnostics csm.dsh csm.gui.dcem csm.gui.websm csm.server csm.hams
- o Refer to [CSM Level Check](#)
- o lspp -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
refer to [AIX52 RSCT level check](#)
- o refer to [AIX53 RSCT level check](#)

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

Assign IP Addresses for SNIs

"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 \(SA38-0646-00\)](#) 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 (see step 10c below).

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.

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-03.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:

- `oslevel -r`
- `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.

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 given in [Table 2](#)" HPS/SNI Detailed LP Level Check"

Notes:

i. 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.

ii. 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

- i. Check which kernel is currently in use: `bootinfo -K` a response of "32" is a 32bit Kernel
- ii. `ln -fs /usr/lib/boot/unix_64 /unix`
- iii. `ln -fs /usr/lib/boot/unix_64 /usr/lib/boot/unix`
- iv. 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
- lspv -l hdisk0 |grep hd5
 - hd5 1 1 01..00..00..00 N/A (hdisk0 is your ipldevice)
- v. Issue: bosboot -ad /dev/<ipldevice> (eg. bosboot -ad /dev/hdisk0)
- vi. Reboot: shutdown -Fr
- vii. 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)Would generate this response:

```
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`

Dynamic Host Configuration Protocol (DHCP)

- Disable Dynamic Host Configuration Protocol (DHCP) on the HMC and reboot the HMC.
- 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](#)
- 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`
 - `openssh.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.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`
 - [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 successfully completing Steps 6a & 6b above 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:
 - i. On the HMC GUI, in the "HMC Management" panel, select the "HMC Configuration" option.
 - ii. Select the "Customize Network Settings" option.
 - iii. Under the "LAN Adapters" tab, select the service network interface for your HMC, example, "Ethernet eth0" and click on "Details"
 - iv. Under the "LAN Adapter" tab, in the "Local Area Network Information sub-section, select the "open" option.
 - v. 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.
 - vi. Click *OK*.
 - vii. The HMC will have to be rebooted in order for the changes to be in effect.

Power off CECs

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

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.

- i. Install the HPSNM_Serv_Pack.zip on the HMC and apply the csm.hpsnm PTF fileset to the CSM MS.
- ii. 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](#):

- [230](#)
- [235](#)
- [240](#)

For the CSM MS:

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

CSM Packages for AIX and Linux are available at:

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

Installation Instructions:

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
- b. 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.

IBM HPS POWER5 Readme - Service Pack 12

[Introduction](#)

[Supporting documentation](#)

[Component updates](#)

[Installation sequence](#)

[AIX LP](#)

[Known problems and workarounds](#)

[Fix list](#)

[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. They are available from [pSeries support fixes](#).

Order by individual APAR OR Order ALL Fixes for this HPS POWER5 Service Pack: [APAR IY92413](#).

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 on the 'pSeries support fixes' website. There you will be presented with an opportunity to get the latest available fixes for that component.

Alternatively, clicking on the 'Order All the Latest' link will open a new page on the 'pSeries support fixes' website. There you will be presented with an opportunity to get all the latest available fixes for this Service Pack.

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:

```
islpp -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	2.4.3	2.3.3
HPS/SNI (CSS)	1.2.0	1.1.3
PE	4.3.0	4.2.2
LoadL	3.4.0	3.3.2
GPFS	3.1.0	2.3.0
CSM	1.5.1	1.5.1
RSCT	2.4.6	2.3.10
ESSL	4.2.0	4.2.0
PESSL	3.2.0	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 [pSeries support fixes](#).

LP	Release	Component ID	APAR	Level check
VSD	410	5765G2602	IY92233	rsct.vsd.cmds 4.1.0.16 rsct.vsd.rvsd 4.1.0.14 rsct.vsd.vsdd 4.1.0.18 rsct.vsd.vsdrm 4.1.0.7
LAPI AIX52 only	233	5765G2601	IY84756	rsct.lapi.nam 2.3.3.2 rsct.lapi.rte 3.3.3.7 rsct.lapi.samp 2.3.3.0
LAPI AIX53 only	243	5765G2601	IY92234	rsct.lapi.rte 2.4.3.1 rsct.lapi.samp 2.4.3.0 rsct.lapi.nam 2.4.3.0
HPS/SNI AIX52 only	113	5765G2400	IY92235	devices.chrp.IBM.HPS.rte 1.1.3.10 devices.common.IBM.sni.rte 1.1.3.8 devices.common.IBM.sni.ml 1.1.3.2 devices.common.IBM.sni.ntbl 1.1.3.1 devices.chrp.IBM.HPS.hpsfu 1.1.3.0
HPS/SNI AIX53 only	120	5765G2400	IY92236	devices.chrp.IBM.HPS.rte 1.2.0.6 devices.common.IBM.sni.ml 1.2.0.2 devices.common.IBM.sni.rte 1.2.0.6 devices.msg.en_US.chrp.IBM.HPS.rte 1.2.0.2 devices.chrp.IBM.HPS.hpsfu 1.2.0.0
PPE	422	5765F8300	IY92410	ppe.poe 4.2.2.6 ppe.man 4.2.2.2 ppe.perf 4.2.2.2 ppe.pvt 4.2.2.0 ppe.loc.license 4.2.0.0 ppe.dpcl 3.3.3.0
PPE	430	5765F8300	IY92237	ppe.poe 4.3.0.1 ppe.man 4.3.0.0 ppe.perf 4.3.0.1 ppe.pvt 4.3.0.0 ppe.loc.license 4.3.0.0 ppe.dpcl 3.4.2.0
LoadL	332	5765E6900	IY92411	LoadL.full 3.3.2.7 LoadL.so 3.3.2.7 LoadL.tguides 3.3.1.0 LoadL.loc.license 3.4.0.0

LoadL	340	5765E6900	IY92238	LoadL.full 3.4.0.1 LoadL.so 3.4.0.1 LoadL.webui 3.4.0.1 LoadL.tguides 3.4.0.0 LoadL.loc.license 3.4.0.0
GPFS	230	5765F64AP	IY92239	gpfs.base 2.3.0.18 gpfs.docs.data 2.3.0.8 gpfs.msg.en_US 2.3.0.12
GPFS	310	5765G66P	IY92240	gpfs.base 3.1.0.7 gpfs.docs.data 3.1.0.2 gpfs.msg.en_US 3.1.0.6
CSM	151	5765E88AP	IY92412	csm.server 1.5.1.3 csm.client 1.5.1.2 csm.core 1.5.1.3 csm.deploy 1.5.1.2 csm.dsh 1.5.1.3 csm.bluegene 1.5.1.1 csm.diagnostics 1.5.1.0 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

RPMS:

```

expect 5.32-1
openCIMOM 0.8-1
openssl 0.9.7d-2
AIX-rpm 5.2.0.40-1
tcl 8.3.3-1
tk 8.3.3-1
conserver 8.1.7-2

```

RSCT	2.3.10	5765F07AP	IY92243	rsct.opt.storagerm 2.3.10.1
				rsct.opt.saf.amf 2.3.10.1
AIX52				rsct.exp.cimrm 2.3.10.0
only				rsct.crypt.rsa512 2.3.10.0
				rsct.crypt.rsa1024 2.3.10.0
				rsct.crypt.des 2.3.10.0
				rsct.crypt.aes256 2.3.10.0
				rsct.crypt.aes128 2.3.10.0
				rsct.crypt.3des 2.3.10.0
				rsct.core.rmc 2.3.10.1
				rsct.core.sr 2.3.10.0
				rsct.core.errm 2.3.10.0
				rsct.core.auditrm 2.3.10.0
				rsct.core.fsrms 2.3.10.0
				rsct.core.hostrm 2.3.10.1
				rsct.core.utils 2.3.10.1
				rsct.core.gui 2.3.10.0
				rsct.core.sec 2.3.10.1
				rsct.core.sensorm 2.3.10.1
				rsct.core.lprm 2.3.10.0
				rsct.compat.clients.rte 2.3.10.0
				rsct.compat.clients.sp 2.3.10.0
				rsct.compat.clients.hacmp 2.3.10.0
				rsct.compat.basic.rte 2.3.10.0
				rsct.compat.basic.sp 2.3.10.0
				rsct.compat.basic.hacmp 2.3.10.0
				rsct.basic.rte 2.3.10.1
				rsct.basic.sp 2.3.10.0
				rsct.basic.hacmp 2.3.10.0
RSCT	2.4.6	5765F07AP	IY92244	rsct.opt.storagerm 2.4.6.1
				rsct.opt.saf.amf 2.4.6.1
AIX53				rsct.exp.cimrm 2.4.6.0
only				rsct.crypt.rsa512 2.4.5.0
				rsct.crypt.rsa1024 2.4.5.0
				rsct.crypt.des 2.4.5.0
				rsct.crypt.aes256 2.4.5.0
				rsct.crypt.aes128 2.4.5.0
				rsct.crypt.3des 2.4.5.0
				rsct.core.rmc 2.4.6.1
				rsct.core.sr 2.4.6.0
				rsct.core.errm 2.4.6.0
				rsct.core.auditrm 2.4.6.0
				rsct.core.fsrms 2.4.6.0


```

rsct.core.hostrm 2.4.6.1
rsct.core.utils 2.4.6.1
rsct.core.gui 2.4.6.0
rsct.core.sec 2.4.6.1
rsct.core.sensorrm 2.4.6.1
rsct.core.lprm 2.4.6.0
rsct.compat.clients.rte 2.4.6.0
rsct.compat.clients.sp 2.4.6.0
rsct.compat.clients.hacmp 2.4.6.0
rsct.compat.basic.rte 2.4.6.0
rsct.compat.basic.sp 2.4.6.0
rsct.compat.basic.hacmp 2.4.6.0
rsct.basic.rte 2.4.6.1
rsct.basic.sp 2.4.6.0
rsct.basic.hacmp 2.4.6.0

```

Parallel	320	5765F8400	PK21664	<pre> 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 </pre>
ESSL				
Parallel	330	5765F8400		<pre> pessl.rte.common 3.3.0.0 pessl.rte.hv 3.3.0.0 pessl.rte.rs1 3.3.0.0 pessl.rte.smp 3.3.0.0 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 </pre>
ESSL				
ESSL	420	5765F8200	PK02186	<pre> essl.rte.common 4.2.0.4 essl.rte.rs1 4.2.0.4 essl.rte.rs2 4.2.0.0 essl.rte.smp 4.2.0.4 essl.rte.mp 4.2.0.0 essl.rte.up 4.2.0.0 essl.man.en_US 4.2.0.0 essl.loc.license 4.2.0.0 </pre>

IBM HPS POWER5 Readme - Service Pack 12

[Introduction](#)[Supporting documentation](#)[Component updates](#)[Installation sequence](#)[AIX LP](#)[Known problems and workarounds](#)[Fix list](#)[Readme archive](#)

Restrictions/Known Problems/Workarounds

This pSeries HPS POWER5 release does not support concurrent firmware installation in an operational HPS POWER5 cluster pending further testing of this function.

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 reenable 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 reenable 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

[LAPI hang](#)[No support for redundant cluster service networks](#)[HPS/SNI Protocol stack requires all sni ... in all servers to be on the same network ...](#)[When updating Licensed Internal Code, if the "High Performance Switches" ...](#)[No support in Service Focal Point for collecting VPD for High Performance Switches](#)[B1818600 error is posted from Perc](#)[Service Focal Point \(SFP\) flags a switch board or riser with a bad BIST and/or bad Signature.](#)[The HPS Network Manager End-Point View and lsswendpt command show cage 16 as cage 0.](#)[VPD databases lost on the HMC](#)[When verifying switch chip ports, you should look for Down:No Signal.](#)

[Diagnostic SRC D103B0FF in error logs](#)[Cluster Ready Hardware Server data may not be restored after an upgrade install of the HMC ...](#)

[Retain Tip Abstract:](#)

When a LAPI hang is reported that does not involve a timeout and the new LAPI_Xfer types LAPI_HW_PUT or LAPI_HW_GET are used then it is possible that this problem has been reproduced. This problem is restricted to jobs running in user space mode using the HPS switch adapter.

Problem Description:

When a LAPI program uses the new LAPI interfaces for hardware get or put using the LAPI call LAPI_Xfer with the LAPI_HW_PUT or LAPI_HW_GET type and request target notification, it is possible that the program will hang. It is unlikely that this problem will be hit on a system with all links known to be reliable since it depends on the last packet of an RDMA operation being only partially completed. This problem is restricted to jobs running in user space mode using the HPS switch adapter.

Customer Impact:

Customers starting development of new programs using this new LAPI interface may have the program suspend operation because RDMA message completion notification never arrives.

Workaround: Download the following files

1SF235_181_160.rpm

01SF235_181_160.xml

[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:
 - i) 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.
 - ii) 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:
 - i) `mv /usr/sni/aix52/cfgsni /usr/sni/aix52/cfgsni.orig`
 - ii) Reboot server
 - iii) For each snix on the server: `rmdev -d -l snix`
 - iv) `mv /usr/sni/aix52/cfgsni.orig /usr/sni/aix52/cfgsni`
 - v) `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:

- i) mv /usr/sni/aix52/cfgsni /usr/sni/aix52/cfgsni.orig
- ii) Reboot server.
- iii) For each snix on the server: rmdev -d -l sniX
- iv) mv /usr/sni/aix52/cfgsni.orig /usr/sni/aix52/cfgsni
- v) cfgmgr -s

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.

[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.

[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.

[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.

[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

[Retain Tip Abstract:](#) VPD databases lost on the HMC

Search Keys:

Trace, Tracelogd, VPD, ELA

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

[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 unuseable 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".

[Retain Tip Abstract:](#) Diagnostic SRC D103B0FF in error logs

Search Keys:

D103B0FF 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

[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:

- log in as root
- check if the directory `"/var/ct.backup"` exists with a timestamp of when the save data was performed.
- If it exists then - run the command

```
"/usr/sbin/rsct/bin/ctrestore"
```

- wait a few minutes for the ClusterPeerDomain to be restored and activated.
- 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:
- mount `/mnt/upgrade`
- verify the `/mnt/upgrade/RSCTSaveUpgrade.tar` file exists.
- If the tar file exists then run:

```
tar -xvf /mnt/upgrade/RSCTSaveUpgrade.tar
```

- check if the directory `"/var/ct.backup"` exists.
- If it exists then run the command

```
"/usr/sbin/rsct/bin/ctrestore"
```

- wait a few minutes for the ClusterPeerDomain to be restored and activated.
- 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.

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.

IBM HPS POWER5 Readme - Service Pack 12

[Introduction](#)

[Supporting documentation](#)

[Component updates](#)

[Installation sequence](#)

[AIX LP](#)

[Known problems and workarounds](#)

[Fix list](#)

[Readme archive](#)

HPS Service Pack Abstract Fix List - By Component

<input type="checkbox"/> HMC	<input type="checkbox"/> VSD	<input type="checkbox"/> Parallel Environment	<input type="checkbox"/> CSM
<input type="checkbox"/> SNM	<input type="checkbox"/> LoadLeveler	<input type="checkbox"/> RSCT	
<input type="checkbox"/> GFW	<input type="checkbox"/> HPS/SNI (Devices)	<input type="checkbox"/> GPFS	

HMC

Refer to the HMC web site for specific details:

<http://www14.software.ibm.com/webapp/set2/sas/f/hmc/home.html>

SNM

1.4.1.16

ELA BAD LOCATION CODE

HPS SWITCH BOARD POWER CYCLE CAN CAUSE PACKET LOSS

RUNNING HALT -Q WHEN STARTING FNM - INITTAB CORRUPTION

fnmd can start before hardware server is ready

Remove SmaGetLinkInfo() from early HW discovery<

close 2nd hole in SMALINKUP flow

HPS Switch TOD Enhancements

HPS Switch Tod Enhancements

No TOD backup and TOD failover with SqH+

FNM core dump on c50mgrs during CEC slow boot

FNM core dump with PTF16 respin#3

Allow P5 H/H+ SMAs to be master (but not backup)

GFW

Refer to the Global Firmware web site for specific details: <http://www14.software.ibm.com/webapp/set2/firmware/gjsn>

VSD

4.1.0.21

Invalid recv message handling
Device Support Enhancements
vsd hangup

4.1.0.22

lower RcvSigRdyFlg retry time
System crash in RcvRmt
system crash in RcvCtl

HPS/SNI

1.1.3.10

P5 NODES CRASH WHEN SNI_DD RUNS OUT OF LARGE PAGES
Remove ifdef around open retry loop
ERRPT MSG 284BDA76 INCORRECTLY LOGGED

1.2.0.6

P5 NODES CRASH WHEN SNI_DD RUNS OUT OF LARGE PAGES
Remove ifdef around open retry loop
ERRPT MSG 284BDA76 INCORRECTLY LOGGED

Parallel Environment

4.2.2.6

DOCUMENTATION NEEDED ON DCPLS AFFECT ON 64-BIT FORTRAN CODE
INCORRECT MESSAGE SIZE RETURNED FROM MPI_PROBE
MEMORY LEAK IN MPI COMPONENT
THE STORAGE USED FOR DERIVED DATA TYPE IS NOT FREED

4.2.2.7

PCT HANGS AT THE END OF THE EXECUTION
Retrofit fix for large page warning/kill capability
PREMATURE PULSE TIMEOUTS WITH REDIRECTED STDIN AND CHECKPOINT

BUG IN LAPI_DGSM_GOSUB COMMAND OF DGSP GENERATED BY MPI_GATHER
 Miscellaneous NTBL Sample flag cleanup
 File /tmp/s.pedb.. has permission 777
 build failed due to 139190 Undefined symbol
 MPCPI LOCK LEAK PROBLEM

4.3.0.1

poe.limits file does not support new MP_BUFFER_MEM syntax
 Jobs hang in lapi_shm_init on SLES 9
 PE Benchmark incorrect MPI trace clock synchronization
 PE Benchmark usermarker error with 64 bit app/64 bit kernel
 Improve the interaction between adapter interrupts and no-poll
 pdb fails if MP_HOSTFILE=NULL or is not set
 13C: MPI_Bcast tuning needed
 0032-117 User pack recv buffer too small
 Linux: update PKD_DATA for PE/LAPI to show correct VRMF
 BUG IN LAPI_DGSM_GOSUB COMMAND OF DGSP GENERATED BY MPI_GATHER
 Linux: pe_install.sh failed to upgrade ppe rpms from GA to PTF1
 Linux: Profiling (gprof) timer expired on PowerPC
 Linux:PE interactive job can't submit to LL if msg_api=mpi,lapi
 Miscellaneous NTBL Sample flag cleanup

Loadleveler

3.3.2.6

LLQ -F AND LLQ -R SHOULD RETURN FAVORED_JOB STATUS AS AN OPTIO
 Ability to query LL version in MultiCluster env
 llclass -X broken between cs3c2 and mercury
 job remains in ep too long
 ON BLUE GENE THE LLCANCEL CAN TAKE A LONG TIME
 THE LLMODIFY -W DOES NOT WORK PROPERLY WITH THE BACKFILL SCHED
 Incorrect casting of variables in LoadL_schedd
 LoadLeveler for Linux: Wrong secondary groups used at run time
 NEGOTIATOR DEADLOCK SCHEDULING SP2 SWITCH
 INC MSG FROM LLQ -S WHEN NOT ENOUGH INITIATORS AND NEED ADAPTE
 Creating Reservation causes Negotiator to crash
 LL MAIL NOT DELIVERED DUE TO INVALID MAIL ADDRESS
 PROBLEM RUNNING A 3.3.1.3 AND 3.3.2.4 MIXED CLUSTER
 llq and llqres display incorrect header and separator
 LoadL PTF needs to show correct level for ptf set 6

3.3.2.7

If i ifconfig and adapter down after preempting a job that was A LOADL RECYCLE ON THE CENTRAL MANAGER WHEN A RESERVATION IS MEMORY LEAK IN SCH_API_EXT.C SAMPLE PROGRAM WHEN EXECUTED LOADL JOBS AT SAME TIME CAN RESULT IN NON-UNIQUE TEMP NAME USER PROLOG ERROR MSG WHEN FAILURE TO SETUP USER ENVIRONMENT LOADLEVELER DOES NOT TRY RECONNECTING TO RSCT- RMC_START_SESSI LOAD LEVELER LOGFILE BEING CREATED WITH OTHER WRITE PERMISSION USING LLMOVEJOB THE REMOTE FILTER IS NOT CALLED BUT LOCAL IS. Bad Function Call Parameter reported by ZeroFault ADAPTER COEXISTENCE BUG SUBMITTING A JOB FROM LL 3.3.1 TO 3.3. Submit job on AIX will failed due to reservation error. Level Bump Defect for sprelmer/ptf7

GPFS

3.1.0.5

32 file systems needs consistency on aix and linux
 DMAPI should limit the number of sessions that can be created.
 mmdeldisk with daemons down fails "twice"
 Linux cannot read file in datashipping mode
 assert:oldVal==SlimFatBit OR oldVal==WideFatBit
 Long waiters on c154n06 attempting recovery after memory assert
 long waiters waiting for SG cleanup
 MC:If cipherList AUTHONLY is not defined on a cluster mount
 assert:err == E_OK, file fencing.C, line 974
 Commands: filter CR and last NL
 HA-NFS: Take over IP address
 HA-NFS: Unlink
 bigfs:mmfsck shows negative numbers
 POLICY_RESTORE poolid
 Node in arbitrating state: ProbeClusterThread on ThMutex (TcpRe
 bigfs:sig11 during mmfsck:FsockInodeMap::testInode
 HA-NFS: recover faster
 GPFS thread with tid -1 is no longer running
 mmlsmount fails with "Arg list too long" with > 32 file systems
 HA-NFS: stay fuzzy sequential
 mmsanrepairfs with snapshot
 Assert exp(! didSGMgrBeginUse) in sgmdata.C
 DCR: MULTINODE FCNTL FAIRNESS

waiters: disklease waiting for SG cleanup; EndUse waiting for v
 Histogram bucket request strings
 lost data in known_cluster files
 DEADLOCK MAKING CONNECTIONS
 SEGV SharedHashTabTraversal::nextTokenP
 dump all assert on linux each time traces are cut
 Allow 64 filesystems
 tsinode: remove default fileset and pool output
 DEADLOCK IN NSDCLOSE HPSS:dm_make_handle on Linux returns bad handle
 mmpmon enhancements
 HA_NFS: avoid failing the NFS client
 HPSS:DM_EVENT_REMOVE event now passes file handle being removed
 gpfs.snap has problems w/ssl directory
 Trace and dump token modes as readable
 AIX packaging changes, APAR number (PMR 16200,49R,000)
 mmremote use GPFS_rshPath and/or GPFS_rcpPath, errmsg cleanup
 PTF5 updates
 Bad return code from kxAttachSharedMemory
 Add missing mmchconfig parameters
 Fix inittab startup (544598)
 gpfs_gpl.h and libgpfs_gpl.so Samba share API
 make World fails
 mmedquota message could be more specific
 Linux kernext usecount check
 panic: kxSendFlock
 ASSERTION FAILED: OFP->MNODESTATUSIS(NOT_MN OR BECOMING_MN)
 Panic: kx.C:1244:0:0:05B137F8::fcntlOpsRevokeLockP
 Prefetch buffer queue was not being correctly drained
 Panic: kx.C:1244:0:0:05B137F8::fcntlOpsRevokeLockP
 Panic: SFSUnmountFS__FP11StripeGroupUi HandleMBUnmount__FP14MBU
 fcntl fairness: kxRetryRecover during unmount
 566221 broke Linux force unmount
 Hang in quiesce system and vnop rele

3.1.0.6

design: method to avoid old site.mcr compilation
 mmlssnapshot inconsistencies with non-root users
 MMFSCK UNABLE TO FIX FILESYSTEM WHEN DIRECTORY BLOCK SPLIT FAI
 NEEDED: A METHOD TO REQUEST GPFS EXPAND THE INODES PROACTIVELY
 Assert failed during format of new filesystem
 SIGUSR1 handler exits too early
 Realloc moving blocks unnecessarily

assert:fsDA.getNValidAddrs==0ORfsDA.getDiskNum!=diskIndex alloc
posix locks interface changed Linux 2.6.17
Changes for 2.6.16
mmlsfileset takes too long
Support for restore policies to POLICY_RESTORE_RULES
Combined shareRelease/shareAcquire does not acquire
ASSERT FAILED: IP == INODEP IN CACHE.C
New linux support
remote mount fails with error 10- no child processes
HA-NFS: Mount only exported file systems
DEADLOCK WAITING FOR FREESPACE RECOVERY
remove TEMP_SHIP_DEFINES
opsdeclare
mmdf -F output
SIGNAL 11 IN HISTOGRAM::SORT
bigfs:mmfsck in infinite loop
Enable POLICY_RESTORE
Signal 11 (SGAllocMap::sendRequestRegion()) when fs is mounted
tsunmount -C all_remote
HA-NFS: Cleanup
mmchfs -Q requires unmount
Handle kxWaitForFlock failures
Remove the check for SLES8 and SLES9 for ppc64
mmchfs mount check
Disallow tsunmount for remote file systems
MMFSCK DIRECTORY REPAIR FIXES
assert:addr.isReservedORaddr.getClusterIdx==clusterIdx cfgmgr.h
mmount and mmumount, all_local and all_remote
mmdelsnapshot hung 'waiting for RPC replies' for sgmMsgQuiesce0
panicked: Z9gpfsFcntlPvP13gpfsVfsData_tS_S_S_P9cxiNode_tXP10cxi
HANG IN DOREPLICATEDREAD
SEGV WHILE RUNNING FORTRAN PROGRAM ON IA64 ODE BUILD
Improve delete performance in snapshots
Default setting for setCtimeOnAttrChange to false
GROUPPROTOCOLDRIVERTHREAD HUNG WAITING FOR REPLIES
assert: !(((lockWordCopy) & 0xFFFFFFFFFFFFFFFFC) dSynch.C, line
mmfsck option parsing
ASSERT: PARENTLEVEL>=1 AND PARENTLEVEL<=INODE.INDIRECTIONLEVEL
mmlsmount of remote filesystem hangs
UNEXPECTED ERROR FROM GPFS_FPUTATTRS, ERRNO = 22
PIT validate work request, fileset scan: free hash table
FORCE UNMOUNT WITH RC=12 ENOMEM ON FSMGR FAILOVER
PTF6 updates
PPC64: mmfslinux' : -1 Unknown symbol in module

gpfs_gpl.h copyright

Junction path did not populate to non fsmgr nodes

Increase concurrency on fast read/write path

Linux tail command differences

MMSHUTDOWN ERROR MESSAGE IN SYSTEM CONSOLE DURING STARTUP

Error while installing gpfs 3.1 PTF6 on BGL ionodes

3.1.0.7

waiter: sgmMsgFsetOps - In kernel waiting to quiesce fileset
DEADLOCK IN GPFS CODE

64fs:mmismount output different when fs table is full

panic:ctlMutexP->lockWord lockGetattr__FP15KernelOperationP13gp

stuck in arbitrating: waiting for LX lock - quotaMsgPrefetchSha

assert: RecoverFileInode bad objectVersion LogRecovery::recover

assert:err==E_OKORerr==E_STALEORerr==E_DAEMON_DEATH svfs.C 405

assert:e==E_OK&&err==E_OK==rctx.replied tscomm.C:2297

mknod broken on 2.6 kernel

dm_path_to_handle with symlinks

PPC64: mmfsenv Segmentation fault

ALLOW SEQDISCARDTHRESHHOLD AND WRITEBEHINDTHRESHHOLD LARGER 1G

DEADLOCK MULT NODES FILEMSGSYNCFILE WAITER

mmcheckquota - sig11: strlen vsnprintf consprintf QuotaMgr::Off

Show xattr types and dm region in tsinode

RESPOND ETIMEDOUT IF HSM RECALL AND O_NONBLOCK ON LINUX

bigfs:negative number in mmfsck output

Traces for inode expansion

Parallelize dm_get_dirattrs (needs DMAPI_PREFETCH_DIRATTRS)

mmchfs -Q manpage

AIX NFS FCNTL LOCKS HUNG

mmmout: mount point

gpfs_fgetattr description

gpfs_f_llseek treating returned offset like error

Linux install had bad manpage names

Sigsegv if openfile referenced before seg mapped in kernel

HANG EARLY IN MMDELDISK

Catch inconsistent alloc region owners

gpfs did not come up on BGL ionides with 3.1 PTF6

Add EXTERNAL POOL rule for HPSS

Mmfsck should not check filesets if first pass skipped

rpm sles 9 insserv: lustrefs facility \$local_fs, skipped

inconsistent behavior of dm_set_fileattr() on Linux GPFS

CANNOT COMPILE GPFS SAMPLES UTILITIES ON LINUX PPC64

cfgmgr.C:9808: void logAssertFailed - TscReply handler_error
ACCESS NOT BLOCKED WHEN FILE CACHED
STALE NFS FILE HANDLE MESSAGES ON LINUX 2.4
NFS shrink-to-fit logic at fsync
Internal remount functionality
Dump buffers for single inode
Build problem getSiteMcrVer
Smart pointer assert in SFSSetXAttr
Read and write code path perf
negative number after doing an mmchfs -F
mmmount use of tsremount
fcntl fairness: Handle GETLK and non-blocking SETLK
cleanup
locking fairness problem
Designated metanode relinquishing
minus K appears twice in mmchfs usage message
DEADLOCK WAITING FOR MMDEVDISCOVER TO FINISH
mmcrfs needs -j in usage
KSH: mmcrfileset:line 160: syntax error at line 168
tsremount cannot handle the /dev prefix
Remove delvsd references
Convert the ACL header fields for GPFS_GETACL_STRUCT
kernel BUG at fs/inode.c:253!
Fix AIX dump kthreads
Clean stale buffers performance
NFSv4 ACL directory permission check
Handle pagepool errors better
PTF7 updates
Kernel open source
Build break in ODE
Avoid kernel tainted bit on 2.4 Linux kernel
WRONG ERROR RETURN FROM POLL
Spinlock hang in trace
Wrong UTS_RELEASE used in build
128K blocksize
Remove a PCE_26_PCE special case
Debug tools
Enable DMAPI_PREFETCH_DIRATTRS
No gpfs.src rpm available for test
Rename DMREGION to COMANAGED
NFS4 ACLS NOT ALLOWING SOME CREATES IN DIRECTORIES
Move gpfs_prealloc to libgpfs_gpl
kx/linux open source code missed
WAITING FOR LX LOCK ON ROOT DIRECTORY

gpfs.src rpm has incorrect version
opensrc make World fails
open source: bypass kxCheckBuildTime if not CHECK_KXTIMESTAMP
REMOTE FS NSD IO ERRORS AFTER NSD SERVER CHANGE
mmpmon service_type changed
Do not change license information to GPL yet
ASSERT FAILED: !(CCP->HASJOINED() && CCP->ISXCLUST(NADDR))
FRAGMENT AND DATABLOCK ERRORS AFTER TRUNC

CSM

1.5.1.3

350-4C:Nodegrp noderange syntax (preceding - problems)
THE STATUS OF AN MS THAT IS ALSO AN MN GOES UNKNOWN AFTER RMC
MM:fld:CSMAgentRM ran out memory after being up for 6 month
updatehwdev -I fails if the HMC devices are defined by Hostname
LX: IBM.DMSRM core dump during application of PTF2 (spin 4)
1350-6A:Two agetty lines in inittab with SLES 9 SP3
RMRHWS CAN FAIL AND GIVE NO INDICATION
1Q8 - bmc2d core dumps when closing console
rfwflash fails to update more than one CEC in a frame
1Q8 - SysRq in serial console not working
UPDATENODE FAILS IF PING PORT IS BLOCKED
Add preinstall support for x3550, x3650, x366
Add Crichton Blade preinstall support
Use of adapters other than en1 for blade installs
rfwflash -a -t power does not work with ML and SF frame
ssh setup failure
/csminstall fills up, but still says OK
Almaden got DB2 error when monitoring Blue Gene
ECMWF: frame10 node1 not seen by HPSNM.
reventlog command not returning SP logs!
Unable to use cd1 when creating resources
Incorrect report when the run failed.
ADDING IMMEDIATELY POWER OFF OPTION ON CSM RPOWER COMMAND
rmpeer does'nt work and returns no err if peerdomain is offline
Update VRMF to 1.5.1.3
UPdate Esopus HA MS with 137093 fix for lsrsrc change
Install httpd failed on RHEL4U3(x86_64) when running csmsetupks
rfis0629a build breaks
CFMUPDATENODE -F FLAG DOESNT WORK RIGHT IF FILE._GROUP EXISTS
bmc2d dumps core when closing multiple console sessions

rmnode failed because /etc/dhcpd.conf does not exist
 DSH SOMETIMES LOSES BLANK LINES IN THE OUTPUT
 Porting:CSM supports RHEL3U8 on x/pSeries
 Update CSM README for 1.5.1.3
 pLinux: setbootdisk not set boot device correctly on some vLpar
 RH: UUID installs fail when gateway or nameservers specified
 CSM prompts for wrong CD names for SLES 9
 HWS should compare MTMs in CRHS only
 rfwflash not working with multiple targets of same type
 Incorrect message function called
 change 1.5.1.1 to 1.5.1.3
 Porting: porting changes

1.5.1.4

CFMUPDATENODE CREATES A DIR WITH BAD PERMS IF ._HOSTNAME EXIST
 Fix dsh/dcp SMIT panels
 RFWSCAN SHOULD UNSET DSH_PATH
 getserviceableevents should delete DSH_PATH
 DMSRM core dump during updatenode
 New BMC lflash package for x346 fails with rfwflash
 rconsole xterm window size wrong if -o option used
 DSH NO LONGER HANDLING MULTIPLE COMMANDS FROM A STANDARD INPUT
 Update VRMF and READMEs for CSM 1.5.1.4
 INSTALLNODE RUNS CFM ON ALL THE NODES AFTER FIRST BOOT

RSCT

2.3.10.1

keep shm size as reasonable limit in large pagesize
 User able to change const Name in harvested IBM.AgFS
 IBM.CONFIGRM USING TOO MUCH CPU DUE TO HAGSGLSM MSG QUEUE ISSU
 nfs filesystem flipflops between online and offline when startd
 CsiAttrs shouldn't be marked as a req attr for SAF_CSI
 LX: Able to change the Name attr of the IBM.VolumeGroup rsrc
 possible memory leak in RMACCommandGroup.C
 HAGSD MIGHT ASSERT WHEN RECEIVING NON-RSCT PACKETS
 forcerpoffline should fail if IW is specified as an operand
 StorageRM core found
 cthags coredumps on Zorro Linux cluster in prm code
 startdms does not work

problem with softdog and critical TSA resources
 CtSec segfault on corrupted THL file
 Bump hmc.lparcmdrm to 1.3 for
 chrsrc-api core dumps when invoked with -I
 hags core dump due to unnamed block in TraceStream.operator
 rmc core dump abort core dump while apply dvo to debs003a
 Do not ship symlink in /usr/lib/ to libct_amf.so
 AGFileSystem not failed offline when pulling cable
 Resources AgFileSystem on same VG only partially online
 RM might die when N-N batch has fewer targets than commands
 More cthags coredump after scaling RPD on Zorro
 Initial OpState for the stopped node should be OFFLINE
 rsct daemons can not be started
 chsensor returns an incorrect hint.
 startrsrc doesn't fail when mount dir doesn't exist.
 LX: call cu_pick_thread_stacksize() when creating security thre
 Unable to change ProtectionMode or DeviceLockMode in VG or LV
 Bump VRMF for Dvorak/Dvorak52 PTF 1
 attach shm fail in large page sz
 don't set LD_ASSUME_KERNEL on RHEL5
 call cu_pick_thread_stacksize for proper adjustment
 CT: ctmskf core dump at clic.CLiC_dispose
 CRunble: adjust thread stack size to reasonable amount
 FIX TO PHOENIX.SNAP CODE UNDER SNAP -E.
 Incorrect SysMountPoint is harvested
 Dead Man Switch Disabled displays on console when node online
 hang in ct2prm functions
 deadlock in RMRccp::enumerateRcps()
 Correct the CMVC Prolog in SAFRM mdef file

2.4.6.1

keep shm size as reasonable limit in large pagesize
 User able to change const Name in harvested IBM.AgFS
 IBM.CONFIGRM USING TOO MUCH CPU DUE TO HAGSGLSM MSG QUEUE ISSU
 nfs filesystem flipflops between online and offline when startd
 CsiAttrs shouldn't be marked as a req attr for SAF_CSI
 LX: Able to change the Name attr of the IBM.VolumeGroup rsrc
 possible memory leak in RMACCommandGroup.C
 HAGSD MIGHT ASSERT WHEN RECEIVING NON-RSCT PACKETS
 forcerpoffline should fail if IW is specified as an operand
 StorageRM core found on c108f1n15.
 cthags coredumps on Zorro Linux cluster in prm code

startdms does not work
problem with softdog and critical TSA resources
CtSec segfault on corrupted THL file
chrsrc-api core dumps when invoked with -I
hags core dump due to unnamed block in TraceStream.operator
rmc core dump abort core dump while apply dvo to debs003a
Do not ship symlink in /usr/lib/ to libct_amf.so
AGFileSystem not failed offline when pulling cable
Resources AgFileSystem on same VG only partially online
RM might die when N-N batch has fewer targets than commands
More cthags coredump after scaling RPD on Zorro
Initial OpState for the stopped node should be OFFLINE
rsct daemons can not be started
chsensor returns an incorrect hint.
startsrc doesn't fail when mount dir doesn't exist.
LX: call cu_pick_thread_stacksize() when creating security thre
Unable to change ProtectionMode or DeviceLockMode in VG or LV
Bump VRMF for PTF 1
attach shm fail in large page sz
call cu_pick_thread_stacksize for proper adjustment
ctmsskf core dump at clic.CLiC_dispose
CRunble: adjust thread stack size to reasonable amount
FIX TO PHOENIX.SNAP CODE UNDER SNAP -E.
Incorrect SysMountPoint is harvested
Dead Man Switch Disabled displays on console when node online
hang in ct2prm functions
deadlock in RMRccp::enumerateRcps()
Correct the CMVC Prolog in SAFRM mdef file

IBM HPS POWER5 Readme - Service Pack 12

[Introduction](#)

[Supporting documentation](#)

[Component updates](#)

[Installation sequence](#)

[AIX LP](#)

[Known problems and workarounds](#)

[Fix list](#)

[Readme archive](#)

IBM HPS POWER5 FLASH/Readme Archive

[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](#)

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

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

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