

IBM pSeries HPS POWER5 FLASH/Readme - Service Pack 1

Service Pack 1 is comprised of CSM, RSCT and GPFS updates.

The IBM pSeries^R High Performance Switch (HPS) POWER5TM Release requires updates to the Hardware Management Console (HMC), HPS Network Manager (HPSNM), HPS/Switch Network Interface (SNI) LP, Global Firmware (GFW), Power Subsystem Microcode, Cluster Systems Management (CSM) LP, Reliable Scalable Cluster Technology (RSCT) software, AIX 5LTM Version 5.2 and various supporting AIX^R LPs.

This IBM pSeries^R HPS POWER5TM Release details updates to these components:

[Hardware Management Console \(HMC\)](#) V4R5.0 (Build level 20050519.1) + MH00318 Efix (Required Maintenance Fix for V4R5.0 (04-25-2005))

Power Subsystem Microcode 02BP131_125

Global Firmware (GFW) 01SF230_126_120

HPS/SNI LP 1.1.3.2

CSM LP 1.4.1.2

RSCT LP 2.3.6.2

AIX 5L Version 5.2 bos.mp64 5.2.0.61

[AIX LPs \(HPC\)](#) as listed in "[Detailed LP Level Check](#)"

This document also contains general guidelines for upgrading the components listed in [Section 2: "Component Update/Download Information"](#).

These guidelines are intended to be a supplement to the other IBM documents referred to in this document. We strongly advise that you have the referenced documents available before you begin the upgrade process. A list of referenced documents can be found in [Section 1: Supporting Documentation](#).

The Code Levels listed in [Section 2: "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 [Section 3: "Recommended Installation Sequence"](#) is the standard sequence of installation.

Non-standard sequences and/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 it is more than 75% full, look for the directories that contain the most data. [use 'du /var | sort -n']

- [Supporting Documentation](#)
- [Component Update/Download Information](#)
- [Recommended Installation Sequence](#)
- [Restrictions/Known problems/Workarounds](#)
- [HPS Service Pack fix Abstract by component](#)

Section 1: Supporting Documentation

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

Section 2: Component Update/Download Information

Component	Service pack updates	Download sites
HMC	<p>This release requires that the HMC is installed with HMC V4R5.0. Select "HMC_Recovery_CD_V4R5.0"</p> <p>Note: CD images cannot be downloaded directly from this site. Contact your IBM Sales Representative or Business Partner, and order Hardware Feature Code (MES) 0960 for the initial upgrade CDs.</p> <p>----- This release requires the HMC MH00318 Efix applied on HMC V4R5.0.</p>	<p>http://techsupport.services.ibm.com/server/hmc/power5/fixes/v4r5.html</p> <p>----- http://techsupport.services.ibm.com/server/hmc/power5/fixes/ptf_MH00318.html</p>

Select "Required V4 R5.0 MAintenance eFix (MH00318)"

<ftp://techsupport.services.ibm.com/eserver/pseries/hmc/fixes/MH00318.zip>

Power
Subsystem
Microcode

Fix pack 02BP230_131_125

<http://techsupport.services.ibm.com/server/mdownload/other.html>

Note that Power Subsystem Microcode is available up to 1 week earlier than the mentioned website (see download site). Please contact your IBM Service Support Rep (SSR) for the Power Subsystem Microcode if not available on website.

GFW

Fix pack 01SF230_126_120

<http://techsupport.services.ibm.com/server/mdownload/download.html>

Note that GFW is available 3-4 days earlier than the mentioned website (see download site) through your IBM SSR. Please contact your IBM SSR for the GFW package if not available on website. Refer to this website for detailed download and unpacking procedures: <http://techsupport.services.ibm.com/server/mdownload/downproc.html>

AIX

The recommended AIX service level for AIX 5L Version 5.2 is the 5200-06 maintenance package (Specific Fixes APAR IY67913) plus a AIX5L Version 5.2 PTF (Specific Fixes APAR IY71994).

[pSeries support fixes](#)

openCIMOM (CSM requisite)

CSM 1.4 requires RPM update openCIMOM 0.8(5.2)

[openCIMOM update](#)

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

SSL (CSM requisite)

openssl-0.9.7d-2
AIX-rpm-5.2.0.40-1

<http://www.ibm.com/link/oss.software.ibm.com/redirect.shtml/projects/opensshi>

SSH (CSM
requisite)

openssh-3.8.1p1_52.tar.gz

<http://www.ibm.com/link/oss.software.ibm.com/redirect.shtml/projects/opensshi>

HPSNM

The AIX filesets called csm.hpsnm and csm.server should ONLY be installed on the CSM MS and not on production LPARS

<http://techsupport.services.ibm.com/server/csm/download>

This release requires the HPSNM 1.4.1.1 version.

For HPSNM AIX PTF, csm.hpsnm: APAR IY70554, PTF U803700

<http://techsupport.services.ibm.com/server/hmc/power5/fixes/v4r5.html>

For HPSNM RPM, IBMhsc.HPSNM_logon the HMC: APAR MB01081, PTF MH00299

AIX LPs:

This HPS POWER5 Release updates these components: VSD LAPI SNI PE LoadL GPFS CSM RSCT

If you have any of the following AIX LPs components installed, apply the listed APARs which are needed for this HPS P5 Release.

They are available from the website: [pSeries support fixes](#)

Order by individual APAR

OR

Order ALL Fixes for HPS POWER5 Service Pack 1: IY73172

To verify that the service levels for your LPs are current for this HPS P5 Release, on each logical partition issue:

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

Depending on the LP's you have, your output should match these LP fileset levels:

Detailed LP Level Check

LPP	Release	Component ID	APAR	Level check
VSD	410	5765G2602	IY71747	rsct.vsd.cmds 4.1.0.12 rsct.vsd.rvsd 4.1.0.10 rsct.vsd.vsdd 4.1.0.11 rsct.vsd.vsdrm 4.1.0.4
LAPI	233	5765G2601	IY71750	rsct.lapi.rte 2.3.3.2 rsct.lapi.nam 2.3.3.0 rsct.lapi.samp 2.3.3.0
HPS/SNI	113	5765G2400	IY71758	devices.chrp.IBM.HPS. hpsfu 1.1.3.0 devices.chrp.IBM.HPS. rte 1.1.3.2 devices.common.IBM. sni.ml 1.1.3.0 devices.common.IBM. sni.ntbl 1.1.3.1 devices.common.IBM. sni.rte 1.1.3.2
PPE	420	5765F8300	IY71763	ppe.poe 4.2.0.2 ppe.man 4.2.0.1 ppe.loc.license 4.2.0.0 ppe.dpcl 3.3.3.0 ppe.perf 4.2.0.1 ppe.pvt 4.2.0.0

LoadL	330	5765E6900	IY71765	LoadL.so 3.3.0.2 LoadL.full 3.3.0.2 LoadL.loc.license 3.3.0.0 LoadL.tguides 3.3.0.1
GPFS	230	5765F64AP	IY73171	mmfs.gpfs.rte 2.3.0.4 mmfs.base.cmds 3.6.1.6 mmfs.base.rte 3.6.1.6
CSM	141	5765E88AP	IY73169	csm.client 1.4.1.3 csm.core 1.4.1.3 csm.diagnostics 1.4.1.2 csm.dsh 1.4.1.2 csm.gui.dcem 1.4.1.2 csm.gui.websm 1.4.1.1 csm.hpsnm 1.4.1.1 csm.server 1.4.1.3 pssp.pssp_to_csm 1.1.1.0 <>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	236	5765F07AP	IY73171	rsct.basic.rte 2.3.6.3 rsct.compat.basic.rte 2.3.6.0 rsct.compat.clients.rte 2.3.6.0 rsct.core.auditrm 2.3.6.1 rsct.core.errm 2.3.6.0 rsct.core.fsrn 2.3.6.0 rsct.core.gui 2.3.6.0 rsct.core.hostrm 2.3.6.0 rsct.core.lprm 2.3.6.1 rsct.core.rmc 2.3.6.3 rsct.core.sec 2.3.6.3 rsct.core.sensorm 2.3.6.1 rsct.core.sr 2.3.6.1 rsct.core.utils 2.3.6.3 rsct.exp.cimrm 2.3.6.1

Parallel ESSL	320	5765F8400		<p>pessl.loc.license 3.2.0.0 pessl.man.en_US 3.2.0.0 pessl.rte.mp 3.2.0.0 pessl.rte.up 3.2.0.0 pessl.rte.common 3.2.0.0 pessl.rte.rs1 3.2.0.0 pessl.rte.rs2 3.2.0.0 pessl.rte.smp 3.2.0.0</p>
ESSL	420	5765F8200	PK02186	<p>essl.loc.license 4.2.0.0 essl.man.en_US 4.2.0.0 essl.rte.mp 4.2.0.0 essl.rte.up 4.2.0.0 essl.rte.common 4.2.0.2 essl.rte.rs1 4.2.0.2 essl.rte.rs2 4.2.0.0 essl.rte.smp 4.2.0.2</p>

Section 3: Recommended Installation Sequence :

NOTES TO IBM SSR's and SYSTEM ADMINISTRATORS :

For the hardware and software installation of the HPS Cluster, we recommend that the Service Support Representatives and System Administrators go through the [pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide for IBM eServer p5 servers \(SA38-0646-00\)](#), [CSM for AIX5L V1.4.1 Planning and Installation Guide](#) and [CSM for AIX5L V1.4.1 Administration Guide](#) in order to co-ordinate the necessary installation and configuration tasks.

Chapter 6 ("Installation") of the [pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide for IBM eServer p5 servers \(SA38-0646-00\)](#) provides detailed information on the hardware, firmware and software sequence of installation.

1a

- i) Install Hardware Management Console (HMC) software on all HMCs in the cluster.
- ii) Apply the HMC EFIX MH00318 patch on all HMCs in the cluster.

Documentation Reference:

[Hardware Management Console](#)

[Hardware Management Console V4R5.0 Readme](#)

[Hardware Management Console for pSeries Installation and Operations Guide](#)

[pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide for IBM eServer p5 servers \(SA38-0646-00\)](#)

[Hardware Management Console for pSeries Maintenance Guide](#)

HMC Software Download Information:

In the Fixes and updates section on the [HMC Power5servers](#) page, select the Version 4.5 machine code updates option for Recovery CD package information and HMC installation instructions.

Download the [HMC EFIX MH00318 from the Required V4R5.0 Maintenance eFix](#) website and follow the instructions to apply the efix to the HMCs.

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 Power5servers](#) page.

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:

`lshmc -V` shows:

Version: 4

Release: 5.0

HMC Build level 20050519.1

MH00318: Required Maintenance Fix for V4R5.0 (04-25-2005)

Level Check from HMC GUI:

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

Version: 4

Release: 5.0

HMC Build level 20050519.1

MH00318: Required Maintenance Fix for V4R5.0 (04-25-2005)

Notes:

Have your IBM SSR download the most recent copy of the HPS Install and Planning Guide.

Install the HMC code using the instructions in the HPS Guide, [pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide for IBM eServer p5 servers \(SA38-0646-00\)](#).

Review the HMC information on the web page where you downloaded the images. For new HMC installation(s), follow the instructions as described in Chapter 3 ("Installing and Configuring the HMC") of the [Hardware Management Console for pSeries Installation and Operations Guide](#).

Install Note: 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 imbedded kernel will not run well when this option is enabled. This setting must be disabled before upgrading to HMC R4V5.

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

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 \(SA38-0646-00\)](#).

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

- 1b 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.

1c	<p>Enable the Cluster Ready Hardware Server 8899:tcp port on the service network on the HMCs in the cluster.</p> <p>Directions to enable Cluster Ready Hardware Server port Open HMC GUI, Customize Network Configuration Select Ethernet interface Click on “Details” bottom Click on “Firewall” Highlight Cluster Ready Hardware Server 8899:tcp from upper window Click on “Allow incoming” bottom Check the lower window, verify Cluster Ready Hardware Server appears there.</p> <p>>> Reboot HMC after the configuration.</p>
2	<p>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 (SA38-0646-00) for more details.</p>
3	<p>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 servers (SA38-0646-00) for installation details.</p> <p>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.</p> <p><u>Documentation Reference:</u> pSeries High Performance Switch (HPS) Planning, Installation and Service Guide for IBM eServer p5 servers (SA38-0646-00)</p> <p><u>Power subsystem microcode Download Information:</u> https://techsupport.services.ibm.com/server/mdownload/other.html</p> <p><u>Level Check:</u> 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.</p> <p>EC Number LIC Type Machine Type/Model/Serial Number Installed Level Activated Level</p>

Accepted Level

02BP230 Power Subsystem 9458-100*<SerialNumber> 131 131 131

(Note that the Accepted LIC level is the last successfully installed and activated power subsystem LIC level)

Notes:

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 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 \(SA38-0646-00\)](#) 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 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.

Also, 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 \(SA38-0646-00\)](#)

Managed system firmware (GFW) Download Information:

<http://techsupport.services.ibm.com/server/mdownload2/download.html>

4

Level Check:

To view the **expected** installed and activated managed system 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

01SF230 Managed System 9458-100*<SerialNumber> 126 126 125

(Note that the Accepted LIC level is the last successfully installed and activated managed system LIC level)

Notes:

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

5

Create LPAR definition and assign adapters. Refer to \"Step 1: HMC code load\" 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 \(SA38-0646-00\)](#) for more details.

Install Cluster System Manager (CSM) software on the CSM Management Server(s) in the cluster. The CSM software is part of the base AIX installation.
The recommended AIX service levels for AIX5L version 5.2 is 5200-06 Maintenance package (APAR IY67913) plus AIX5L Version 5.2 PTF (APAR IY71994).

Preliminary Notes:

CSM 1.4 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> are pre-requisites for applying software for csm.server.1.4.1.1 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 V1.4.1 Planning and Installation Guide (SA23-1334-00)

For information on CSM configuration, refer to CSM for AIX5L V1.4.1 Administration Guide (SA23-1334-00)

[pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide for IBM eServer p5 servers \(SA38-0646-00\)](#)

AIX5L version 5.2 Maintenance Package 5200-06 Download Information:

To download:

Go to <http://www-1.ibm.com/servers/eserver/support/pseries/aixfixes.html>

Select Maintenance Packages and choose AIX52

Specify your \"Current level\"

Specify \"Desired level\": 5200-06 and click \"GO\"

Follow instructions on this page for downloading the gzip file.

6a

AIX5L version 5.2 PTF1 (APAR IY71994) Download Information:

To download:

Go to <http://www-912.ibm.com/eserver/support/fixes/search.jsp?system=2&release=5.2>

Search by APAR number or abstract

Specify the APAR number, \"IY71994\"

Follow instructions on this page for downloading the file.

Level Check:

Level check AIX by running the following commands on the CSM Management server:

```
[/]> oslevel -r
```

```
5200-06
```

```
[/]> lspp -L bos.mp
```

```
Fileset Level State Type Description (Uninstaller)
```

```
-----  
bos.mp 5.2.0.61 A F Base Operating System Multiprocessor Runtime
```

Notes:

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.

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

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).

Apply the RSCT and CSM PTF update images for AIX on the CSM Management Server.

RSCT PTF Download Information:

To download, click on [RSCT](#)

Select \"AIX5200-06\" for the Maintenance level and click ok

Click on the \"Continue\" button at the bottom of the page.

Follow instructions for downloading the RSCT filesets.

CSM PTF Download Information:

To download, click on [CSM](#)

Select \"AIX5200-06\" for the Maintenance level and click ok

Click on the \"Continue\" button at the bottom of the page.

Follow instructions for downloading the CSM filesets.

Level Check:

6b

```
[/]> lspp -L csm.client csm.core csm.diagnostics csm.dsh csm.gui.dcem csm.gui.websm csm.server csm.hams
```

Fileset Level State Type Description (Uninstaller)

```

-----
csm.client 1.4.1.2 A F Cluster Systems Management Client
csm.core 1.4.1.2 A F Cluster Systems Management Core
csm.diagnostics 1.4.1.1 A F Cluster Systems Management Probe Manager / Diagnostics
csm.dsh 1.4.1.2 A F Cluster Systems Management Dsh
csm.gui.dcem 1.4.1.2 A F Distributed Command Execution Manager Runtime Environment
csm.gui.websm 1.4.1.1 C F CSM Graphical User Interface
csm.hams 1.4.1.2 A F Cluster Systems Management HA
csm.server 1.4.1.2 A F Cluster Systems Management Server
-----

```

```
[/]> lspp -L rsct.core.sensorrm rsct.core.utils rsct.exp.cimrm rsct.basic.rte rsct.core.lprm rsct.core.rmc rsct.core.sec rsct.basic.hacmp rsct.core.auditrm
```

Fileset Level State Type Description (Uninstaller)

```

-----
rsct.basic.hacmp 2.3.6.1 A F RSCT Basic Function (HACMP/ES Support)
rsct.basic.rte 2.3.6.2 A F RSCT Basic Function
rsct.core.auditrm 2.3.6.1 A F RSCT Audit Log Resource Manager
rsct.core.lprm 2.3.6.1 A F RSCT Least Privilege Resource Manager
rsct.core.rmc 2.3.6.2 A F RSCT Resource Monitoring and Control
rsct.core.sec 2.3.6.2 A F RSCT Security
rsct.core.sensorrm 2.3.6.1 A F RSCT Sensor Resource Manager
rsct.core.utils 2.3.6.2 A F RSCT Utilities
rsct.exp.cimrm 2.3.6.1 A F RSCT CIM Resource Manager
-----

```

7

Define the nodes/LPARs in the cluster. Refer to \"Step 4: Set up the CSM cluster\" 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 \(SA38-0646-00\)](#) for more details.

8

Note: \"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 \"Step 6: Install AIX on the nodes\" 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 \(SA38-0646-00\)](#) for more details.

The recommended service levels of AIX5L version 5.2 is 5200-06 Maintenance package (APAR IY67913) plus AIX5L Version 5.2 PTF (APAR IY71994).

Documentation Reference:

[pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide for IBM eServer p5 servers \(SA38-0646-00\)](#)

For information on AIX installation, refer to [AIX 5L Version 5.2 Installation Guide and Reference \(SC23-4389-03\)](#)

Chapter 12 (\"Installing the AIX operating system on the nodes\") of the [CSM for AIX5L V1.4.1 Planning and Installation Guide \(SA23-1334-00\)](#)

also contains details on installing the AIX operating system on the LPARs.

AIX5L version 5.2 Maintenance Package 5200-06 Download Information:

To download:

Go to <http://www-1.ibm.com/servers/eserver/support/pseries/aixfixes.html>

Select Maintenance Packages and choose AIX52

Specify your \"Current level\"

Specify \"Desired level\": 5200-06 and click \"GO\"

Follow instructions on this page for downloading the gzip file.

AIX5L version 5.2 PTF1 (APAR IY71994) Download Information:

To download:

Go to <http://www-912.ibm.com/eserver/support/fixes/search.jsp?system=2&release=5.2>

Search by APAR number or abstract

Specify the APAR number, \"IY71994\"

Follow instructions on this page for downloading the file.

Level Check:

Level check AIX by running the following command on the logical partitions(LPARs):

```
[/]> oslevel -r
```

```
5200-06
```

```
[/]> lspp -L bos.mp
```

```
Fileset Level State Type Description (Uninstaller)
```

```
-----  
bos.mp 5.2.0.61 A F Base Operating System Multiprocessor Runtime
```

Notes:

The following optional filesets, \"csm.hpsnm.1.4.1.0\" and \"csm.server.1.4.1.0\", packaged with AIX5L version 5.2 should NOT be installed on the LPARs in the cluster.

i) Apply HPS/SNI update images to all LPARs. Note that the installation of AIX52 ML06 in the previous step provides base level 1.1.3.0 images on all LPARs in the cluster.

ii) Reboot all LPARs after a successful update to the latest HPS/SNI levels.

Documentation Reference:

[AIX 5L Version 5.2 Installation Guide and Reference \(SC23-4389-03\)](#)

[Switch Network Interface for eServer pSeries High Performance Switch Guide and Reference](#)

HPS/SNI LP download information:

Download and apply the HPS/SNI PTF update included with the listed APARs required for the HPS POWER5 release from the HPS/SNI location.

Level Check:

Issue any one of the following commands on the LPARs to verify the service level for the HPS/SNI LP's after a successful update:

```
lspp -Lc | egrep \"HPS|sni\" | cut -d : -f 2,3 | sed 's:/ /'
```

or

```
dsh \"lspp -Lc | egrep \\\"HPS|sni\\\" | cut -d : -f 2,3 | sed 's:/ /'|dshbak|more
```

Expected Levels:

devices.chrp.IBM.HPS.rte 1.1.3.2

devices.common.IBM.sni.ntbl 1.1.3.1

devices.common.IBM.sni.rte 1.1.3.2

Notes:

1. 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 to them, at this location.

The High Performance Switch adapter requires TLP usage and these TLP requirements are documented.

A formula to calculate the required TLP is highlighted in Step 6b below.

You should also consult the section on Large page feature on AIX in The [AIX 5L Version 5.2 Performance Management Guide](#)

10a

NOTE: Users need to be aware of the usage of the LoadLeveler^R 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.

2. 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.

NOTE: 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.

To set up your LPARs with 64 bit kernel

1) Check which kernel is currently in use: bootinfo -K

a response of \"32\" is a 32bit Kernel

3) ln -fs /usr/lib/boot/unix_64 /unix

3) ln -fs /usr/lib/boot/unix_64 /usr/lib/boot/unix

4) Determine which rootvg hdisk contains the boot logical volume (usually hd5).

This hdisk will be your \"ipldevice\".

a) lspv |grep rootvg

hdisk0 009b982332a1f9b8 rootvg active

hdisk1 009b982332a2321a rootvg active

b) lspv -l hdisk0 |grep hd5

hd5 1 1 01..00..00..00 N/A

(hdisk0 is your ipldevice)

5) Issue: bosboot -ad /dev/<ipldevice> (eg. bosboot -ad /dev/hdisk0)

6) Reboot: shutdown -Fr

7) Verify 64 bit kernel is running after reboot:

bootinfo -K

64

To set up Large Page Option:

(For configuration details, see Large Page Support in [AIX 5L Version 5.2 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 page used by other applications.

Setup Large Page Option using the vmo command for each node or node group:

```
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.

For 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)

10b 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.

Notes:

Regarding the use of memory by Remote Direct Memory Access (RDMA) and TLP.

Real memory is divided into 2 categories - Small Pages and Large pages.

It is the users responsibility to achieve an optimal balance between the the 2 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 lpgg_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.

Ex:

```
#!/bin/ksh
```

```
real_mem=`lsattr -E -l sys0 -a realmem|awk '{print $2}'`
```

```
lpgg_regions=`vmo -o lpgg_regions|awk '{print $3}'`
```

```
A=$(( real_mem * 1024 ))
```

```
B=$(( lpgg_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 it executing. After it exits the memory is unpinned and freed.

Assign IP addresses to all adapter Switch Network Interfaces(SNIs). Refer to "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 \(SA38-0646-00\)](#) for more details.

10c

Tip: 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
```

- i) Disable Dynamic Host Configuration Protocol (DHCP) on the HMC and reboot the HMC.
- ii) 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 V1.4.1 Administration Guide \(SA23-1334-00\)](#)
- iii) 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 V1.4.1 Administration Guide \(SA23-1334-00\)](#).

Preliminary Notes:

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

Download Information:

- 1. 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
- 2. 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
openssh.base.client 3.8.0.5200 COMMITTED Open Secure Shell Commands
openssh.base.server 3.8.0.5200 COMMITTED Open Secure Shell Server

Documentation Reference:

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 \(SA38-0646-00\)](#).

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 V1.4.1 Administration Guide \(SA23-1334-00\)](#).

[CSM for AIX5L V1.4.1 Planning and Installation Guide \(SA23-1334-00\)](#)

[CSM 1.4 FAQs](#)

Level Check:

1) Check that the right levels of RSCT and CSM update filesets were installed on the CSM Management server in Step 6b above.

```
[/]> lspp -L csm.client csm.core csm.diagnostics csm.dsh csm.gui.dcem csm.gui.websm csm.server csm.hams
```

Fileset Level State Type Description (Uninstaller)

```
-----
csm.client 1.4.1.2 A F Cluster Systems Management Client
csm.core 1.4.1.2 A F Cluster Systems Management Core
csm.diagnostics 1.4.1.1 A F Cluster Systems Management Probe Manager / Diagnostics
csm.dsh 1.4.1.2 A F Cluster Systems Management Dsh
csm.gui.dcem 1.4.1.2 A F Distributed Command Execution Manager Runtime Environment
csm.gui.websm 1.4.1.1 C F CSM Graphical User Interface
csm.hams 1.4.1.2 A F Cluster Systems Management HA
csm.server 1.4.1.1 A F Cluster Systems Management Server
-----
```

```
[/]> lspp -L rsct.core.sensorrm rsct.core.utils rsct.exp.cimrm rsct.basic.rte rsct.core.lprm rsct.core.rmc rsct.core.sec rsct.basic.hacmp rsct.core.auditrm
```

Fileset Level State Type Description (Uninstaller)

```
-----
rsct.basic.hacmp 2.3.6.1 A F RSCT Basic Function (HACMP/ES Support)
rsct.basic.rte 2.3.6.2 A F RSCT Basic Function
rsct.core.auditrm 2.3.6.1 A F RSCT Audit Log Resource Manager
rsct.core.lprm 2.3.6.1 A F RSCT Least Privilege Resource Manager
rsct.core.rmc 2.3.6.2 A F RSCT Resource Monitoring and Control
rsct.core.sec 2.3.6.2 A F RSCT Security
rsct.core.sensorrm 2.3.6.1 A F RSCT Sensor Resource Manager
rsct.core.utils 2.3.6.2 A F RSCT Utilities
rsct.exp.cimrm 2.3.6.1 A F RSCT CIM Resource Manager
-----
```

Notes:

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.

From the [CSM 1.4 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:

- On the HMC GUI, in the `"HMC Management"` panel, select the `"HMC Configuration"` option.
- Select the `"Customize Network Settings"` option.
- Under the `"LAN Adapters"` tab, select the service network interface for your HMC, example, `"Ethernet eth0"` and click on `"Details"`

- Under the "\LAN Adapter\" tab, in the "\Local Area Network Information sub-section, select the "\open\" option.
 - 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.
- Click OK.
- Note: The HMC will have to be rebooted in order for the changes to be in effect.

12 Enusre all CECs in the cluster are in the powered off state before proceeding with Step 13a.

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

i) Install the HPSNM_Serv_Pack.zip and apply the csm.hpsnm 1.4.1.1 PTF fileset on the HMC and CSM MS, respectively.

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 \(SA38-0646-00\)](#).

Download Information:

Download the HPSNM Service Pack 2 package (contains HPSNM_Serv_Pack_2.zip) in the HPSNM updates and fixes" section of [HMC POWER5](#) page.

Download the csm.hpsnm.1.4.1.1 AIX PTF from <https://techsupport.services.ibm.com/server/csm/download> using the APAR IY70554, PTF U803700 details.

Installation Instructions:

Installation on the HMC:

The HPSNM_Serv_Pack_2.zip needs to be installed on all HMCs in the cluster, and each HMC should be rebooted.

Install the HPSNM_Serv_Pack_2.zip file on all HMCs from the HMC GUI using the following options: Licensed Internal Code Maintenance -> HMC Code Update -> Install Corrective Service, then 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 HPSNM_Serv_Pack_2.zip was downloaded to a remote system)

Reboot all HMCs with the HPSNM_Serv_Pack_2.zip installed.

Installation on the CSM Management Server (MS):

Apply the csm.hpsnm.1.4.1.1PTF fileset on the CSM MS.

The csm.hpsnm base fileset is installed on the CSM Management Server as part of the CSM MS installation with the AIX5L 5200-06 Maintenance package (Use the "Automatically increase filesystems" option when doing this installation).

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

More Configuration:

13a

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 \(SA38-0646-00\)](#)

Level Check:

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

(1) Verify that the csm.hpsnm.1.4.1.1 fileset was successfully installed, here is an example:

```
[/]> lsipp -L csm.hpsnm
```

```
Fileset Level State Type Description (Uninstaller)
```

```
-----
```

```
csm.hpsnm 1.4.1.1 A F IBM Switch Network Manager
```

13b

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.

Section 4: Restrictions/Known Problems/Workarounds

Restrictions:

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.

Known Problems/Workarounds:

Retain Tip Abstract: No support for redundant cluster service networks in this IBM pSeries^R HPS POWER5TM Release.

Search Keys: cluster service network, service network, service ethernet

Problem Description: HPS Switch Fabric on Squadrons only supports a single cluster service network.

Customer Impact: No option to support a redundant cluster service network for fault tolerance.

In the event you are having difficulties with the switch fabric and no errors are being logged be sure to check that your Network is operational. In addition verify that you are not logging any errors on a BPC on the A side of a system containing an HPS Switch. If so, please repair.

Workaround: None

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 (be mindful of quotation marks and backslashes) on the CSM MS:

```
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: Power On/Off task targetting an IH node (787D machine type) will not function properly and throw an error code of 74 (0x4a).

Search Keys: 787D, Power On/Off, IH, rc=74

Problem Description: While executing the Service Utility Power On/Off task with a prior CEC target selection of an IH node, the powering on or off of that unit will not function properly. Task was not originally designed to control BPC powered devices, but rather 19" I/O devices.

Customer Impact: To remove power from the IH node, use the EPO switch on the frame.

Workaround: None.

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: HMC not able to retrieve certain logs from FSP

Problem Description: The problem is only an issue when HMC attempts to retrieving informationals, hidden, already ack'd by an HMC or an OS

Customer Impact: Customers may not see certain events/errors get posted on the HMC.

Workaround: Use ASM interface on FSP to see all available logs.

Retain Tip Abstract: SMA Wrap test failure for certain SMA handle passed in by the HMC

Problem Description: Verify SMA Links command is returning invalid SMA Handle or invalid return code (faad). There is a bug in the SMA Handle to SMA Hom object conversion.

Customer Impact: The Verify SMA Links test will not be functional for some SMAs without this fix.

Workaround: Provide test team with patches to get around this defect

Retain Tip Abstract: When running Adapter Wrap Test it must be started from the End-Point view.

Search Keys: Adapter Wrap Test

Problem Description: The Adapter Wrap Test must be started from the End-Point view. If the wrap test is started from the Switch Topology view, and the link selected is an adapter link then the wrong test will be performed. The Switch Management GUI should direct the user to start the test from the End-Point view.

Customer Impact: The wrap test will fail if the adapter test is started from the Switch Topology view. This may cause an adapter to be replaced by mistake, if the adapter is in fact operational.

Workaround: When an adapter needs to be tested, using wrap test, start the test from the End-Point view.

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 IIBM 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

Section 5: HPS Service Pack Fix List Abstract - By Component

GPFS:

ALLOCMAP LAYOUT MAKES SEGMENTS TOO SMALL
LX IA64: ASSERT FAILED: P == WANTP
ERROR IN LSPV WHEN ISSUES AS NON-ROOT
NFS AND SAMBA WRITERS SLOW AFTER LS -L ON ANOTHER NODE
GPFS_ARCH_POWER traceback, PPC cross compiles, tests
MOUNT LOOPING WHEN MANY DISKS WITH LONG NAMES
MMLSFS ON LINUX DOES NOT SHOW -E AND -S VALUES
PROBLEMS OPENING OPENSLL LIBRARY
SIGSEGV PARSING NODE NAMES
MMFS.LOG.LATEST SHOWS RESPAWNING GPFS ACTUALLY WORKING
LINUX PANIC ON MMAPPED EXECUTABLE
KERNEL PANIC AFTER KSFSRETRIEVEDISKADDR ERROR
ASSERT FAILED: DIRECT-COMMON.C LINE 687: FREELEN(FP) == PREVLEN

CSM:

GETADAPTERS COMMAND SKIPS DSH AS DEFAULT METHOD FOR THE PR

RSCT: (AIX)

122825 LX: MMFS DAEMON DIES DUE TO CU_GETTIMEOFDAY - UNDEFINED SYMBOL
122935 RMCD CORES IN OBS_VAR, RM DAEMONS CORE IN ASYNC, 2610-202 IN D
123028 /VAR/CT/<CLUSTER ID>/LOG/CTHATS/REFRESHOUTPUT FILE HAS 666 PER

RSCT: (linux)

122825 LX: MMFS DAEMON DIES - CU_GETTIMEOFDAY - UNDEFINED SYMBOL
123028 /VAR/CT/<CLUSTER ID>/LOG/CTHATS/REFRESHOUTPUT FILE HAS 666 PER

GFW :

To see defects corrected by this GFW fix pack:

[http://w3.austin.ibm.com/afs/austin/projects/gfw/images/Squadrons/b0611b_0520.](http://w3.austin.ibm.com/afs/austin/projects/gfw/images/Squadrons/b0611b_0520.230/01SF230_126_120.html)

[230/01SF230_126_120.html](http://w3.austin.ibm.com/afs/austin/projects/gfw/images/Squadrons/b0611b_0520.230/01SF230_126_120.html)

-