IBM HPS POWER5 - Service Pack 9

IBM POWER5 High Performance Switch (HPS) Service Pack 9 Addendum X

Table of Contents

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.

Notice

- This Service Pack provides support for release SF240_233 / BP240_190
- This release is to support HPS ONLY.
- Issues moving to BASE GA7 that are resolved in Release Level SF240_219
    - 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_156
- **FRAME:** Power Code (BPC code) - Version BP230_155

235 - Service Pack Code Levels

- **HMC:** Version 5.1 + MH00593
  IBMhsc.HPSNM_log-1.4.1.10-1 RPM (part of HPSNM/FNM)
  IBMhsc.NM_Common-1.0.0.0-12
- **CEC:** Global Firmware (GFW) - Version SF235_209
- **FRAME:** Power Code (BPC code) - Version BP235_157

240 - Service Pack Code Levels

- **HMC:** Version: 5 Release: 2.1 + MH00586 + MH00610
  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 SF240_233
- **FRAME:** Power Code (BPC code) - Version BP240_190

Requirements

- **CSM MS:** AIX 5L Version 5.2M, AIX 5L Version 5.3E, HPSNM/FNM - CSM.HPSNM 1.4.1.14, RSCT PTF - 2.3.8, CSM PTF - 1.5.0
- **LPARS:** HPS/CSS PTFs (LPARs), AIX 5L Version 5.2M/**CSS 1.1.3**, **AIX 5L Version 5.3E**/ **CSS 1.2.0**

Supported LP's: as listed in Detailed LP Level Check (page Table 2)

What's New

SP5 introduced support for **GA7 firmware** - products and features supported.

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.

Upgrade Workflow

SP3 introduced a modified install plan intended to reduce the maintainence 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 acivate 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; then

> shutdown the LPARs
> Power off the CECs
> Activate the Power Code or the GFW Code
> Power on The CEC(s)
> activate the LPARs

Re-start the wokload

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 unforseen 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.

Supporting Documentation

1. [High Performance Switch Planning,Installation, and Service for IBM server p5 servers](#)
2. [Hardware Management Console](#)
3. [eServer pSeries and AIX Information Center](#)
4. [Switch Network Interface for eServer pSeries High Performance Switch Guide and Reference](#)
5. [AIX 5.2 documentation](#)
6. [AIX 5L Version 5.2 Installation Guide and Reference](#)
7. [AIX 5L Version 5.2 Performance Management Guide](#)
8. [Reliable Scalable Cluster Technology (RSCT) Library](#)
9. [General Parallel File System (GPFS) Library](#)
10. [Cluster System Management (CSM) Library](#)
11. [IBM Parallel Environment (PE) Library](#)
12. [LoadLeveler for AIX5L and Linux V3.3](#)

Component Update/Download Information

Table 1. Components, updates, and download locations

Component	Service pack updates	Download sites
HMC This release requires that the HMC is installed with code levels that are related to specific firmware.	**230:** HMC V4R5.0. Select HMC Version 4.5 HMC_Update_V4R5.0_1.zip HMC_Update_V4R5.0_2.zip This release requires efix MH00454. You can download the HMC Recovery CD V4R5.0 package as a set of .iso images by using Download Director or you can download the corrective service file directly from the Internet.	[HMC corrective service](#) Under "HMC corrective service support" Select "HMC V4.x" Under "" Select "HMC 4.5" For the HMC ReadMe Scroll to "Update corrective sevrvice" select "view" In the 'ReadMe' column

235:

HMC V5R1
Select "HMC_Recovery_CD_V5R1.0"

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

This release has required efix PTF MH00593 HMC Maintenance Package Version 5, Release 1.2.

This is a Maintenance Fix that must be installed on HMC Version 5 Release 1.0.

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

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, which you can then use to burn your own DVD.

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

[HMC corrective service](#)

- Under "HMC corrective service support"
- Select "HMC V5.x"
- Under ""
- Select "HMC 5.1"

For the HMC ReadMe Scroll to "Update corrective sevrvice", select "view' In the 'ReadMe' column.

240:

HMC V5R2.1 + MH00586 + MH00610 (MH00610 is only required for HMCs managing 9119-590, **9119-595**, 9406-595, 9116-561 or 9117-570 systems with redundant service processor feature)

This release requires that the HMC is installed with HMC V5R2.1. for **SF240_233 / BP240_190**

For this release (**SF240_233 / BP240_190**) use this link:

HMC_Download_and_Installation.html

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

[Power Subsystem Microcode](#)

AND

[System Firmware](#)

230:
  Fix pack BP230_155
  Fix pack SF230_156

235:
  Fix pack BP235_157
  Fix pack SF235_209

240:
  Fix pack BP240_190
  Fix pack SF240_233

For this release (**SF240_233 / BP240_190**) use this link:
[Power5 Firmware ReadMe](#)

Note:

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:

AIX52M 5200-08 APAR IY77270 U806621

bos.mp64: 5.2.0.85

[AIX](#)

AIX53E 5300-04 APAR IY77273 U806622

bos.mp64: 5.3.0.40

Also Apply Critical fixes as appropriate.

To download Service Packs for AIX visit

- [pSeries support fixes](#)
- [AIX52M](#)
- [AIX53E](#)
- [AIX 5L Critical fixes](#)

Note:

Critical fixes for AIX 5L service is 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.](#)

Quick links for AIX fixes.

Notes:

openCIMOM
(CSM requisite)

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

SSL (CSM requisite)

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

- openssh-3.8.1p1_52.tar.gz

OpenSSH-AIX (sourceforge.net)

On the **CSM MS**:

This release requires HPSNM 1.4.1.14 apply HPSNM AIX PTF U805874 , APAR IY76696

Note:

AIX filesets csm.hpsnm and csm.server, which may be included in AIX Maintainence or Technical Level Update Packages, should **ONLY** be installed on the CSM MS and **NOT** on production LPARS.

For the **CSM MS**

- System p support fixes
- eServer/support/fixes
- Hardware Management Console Support for UNIX servers and Midrange servers

HPSNM

On the **HMC**:

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

235 Level:
HMC v 5.1

- Under "HMC corrective service support"
- Select "HMC V<X>"
- Under "HMC Version <X>"
- Select "HMC <X>"
- Select "HPSNM/IBMNM fixes" tab

HPSNM Service Pack 4 MB01237
MH00461

240 Level:
HMC v 5.2
HPSNM Service Pack 6 MB01376
MH00689

AIX Licensed Program

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 the website: System p support fixes.

Order by individual APAR OR Order ALL Fixes for this HPS POWER5 Service Pack: APAR IY86284.

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 'System p 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 'System p 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:

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

Table of AIX52 and AIX53 LPs

LPP		AIX Version	
	AIX52		**AIX53**
VSD	4.1.0		4.1.0
LAPI	2.3.3		2.4.2
HPS/SNI (CSS)	1.1.3		1.2.0
PE	4.2.2		4.2.2
LoadL	3.3.2		3.3.2
GPFS	2.3.0		2.3.0
CSM	1.5		1.5
RSCT	2.3.9		2.4.5
ESSL	4.2.0		4.2.0
PESSL	3.2.0		3.2.0

Detailed LP Level Check - Component Compatiblity

List of which components are compatible with which AIXversion.

Depending on the LP's you have, your output should match the LP fileset levels listed in the Detailed LP Level Check table following Detailed LP Level Check table order the latest APARs are available from System p support fixes

LPP	Release	Component ID	APAR	Level check
VSD	410	5765G2602	IY84755	rsct.vsd.cmds 4.1.0.15 rsct.vsd.rvsd 4.1.0.13 rsct.vsd.vsdd 4.1.0.15 rsct.vsd.vsdrm 4.1.0.7
LAPI **AIX52 only**	233	5765G2601	IY84756	rsct.lapi.nam 2.3.3.2 rsct.lapi.rte 3.3.3.5 rsct.lapi.samp 2.3.3.0
LAPI **AIX53 only**	242	5765G2601	IY84757	rsct.lapi.rte 2.4.2.5 rsct.lapi.samp 2.4.2.0 rsct.lapi.nam 2.4.2.0
HPS/SNI **AIX52 only**	113	5765G2400	IY84758	devices.chrp.IBM.HPS.rte 1.1.3.8 devices.common.IBM.sni.rte 1.1.3.6 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	IY84759	devices.common.IBM.sni.rte 1.2.0.4 devices.common.IBM.sni.ml 1.2.0.2 devices.common.IBM.sni.ntbl 1.2.0.0 devices.chrp.IBM.HPS.rte 1.2.0.4 devices.chrp.IBM.HPS.hpsfu 1.2.0.0
PPE	422	5765F8300	IY84761	ppe.poe 4.2.2.3 ppe.man 4.2.2.2 ppe.perf 4.2.2.0 ppe.pvt 4.2.2.0 ppe.loc.license 4.2.0.0 ppe.dpcl 3.3.3.0
LoadL	332	5765E6900	IY84763	LoadL.full 3.3.2.1 LoadL.so 3.3.2.1 LoadL.tguides 3.3.1.0 LoadL.loc.license 3.3.0.0
GPFS	230	5765F64AP	IY84765	gpfs.base 2.3.0.12 gpfs.docs.data 2.3.0.8 gpfs.msg.en_US 2.3.0.9
GPFS **AIX53 only**	310	5765F64AP	IY84767	gpfs.base 3.1.0.1 gpfs.docs.data 3.1.0.1 gpfs.msg.en_US 3.1.0.1

[CSM]	151	5765E88AP	IY84768

csm.server 1.5.1.1
csm.client 1.5.1.0
csm.core 1.5.1.1
csm.deploy 1.5.1.1
csm.dsh 1.5.1.1
csm.bluegene 1.5.1.0
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 1.5.1.0
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

IY84769

rsct.basic.hacmp 2.3.9.1
rsct.basic.rte 2.3.9.1
rsct.basic.sp 2.3.9.0
rsct.compat.basic.hacmp 2.3.9.0
rsct.compat.basic.rte 2.3.9.0
rsct.compat.basic.sp 2.3.9.0
rsct.compat.clients.hacmp 2.3.9.0
rsct.compat.clients.rte 2.3.9.0
rsct.compat.clients.sp 2.3.9.0
rsct.core.auditrm 2.3.9.0
rsct.core.errm 2.3.9.1
rsct.core.fsrm 2.3.9.0
rsct.core.gui 2.3.9.0
rsct.core.hostrm 2.3.9.0
rsct.core.lprm 2.3.9.0
rsct.core.rmc 2.3.9.1
rsct.core.sec 2.3.9.0
rsct.core.sensorrm 2.3.9.0
rsct.core.sr 2.3.9.0
rsct.core.utils 2.3.9.1
rsct.opt.storagerm 2.3.9.1

[RSCT]

239   5765F07AP

AIX52 only

[RSCT]
 	
AIX53 only
	245	5765F07AP	IY84770	`devices.chrp.base.ServiceRM 1.2.0.90`
				`devices.chrp.base.ServiceRM 1.3.0.45`
				`rsct.core.rmc 2.4.5.1`
				`rsct.core.utils 2.4.5.1`
				`rsct.core.errm 2.4.5.1`
				`rsct.opt.storagerm 2.4.5.1`
				`rsct.basic.rte 2.4.5.1`
				`rsct.basic.hacmp 2.4.5.1`
				`rsct.basic.sp 2.4.5.0`
				`rsct.compat.basic.hacmp 2.4.5.0`
				`rsct.compat.basic.rte 2.4.5.0`
				`rsct.compat.basic.sp 2.4.5.0`
				`rsct.compat.clients.hacmp 2.4.5.0`
				`rsct.compat.clients.rte 2.4.5.0`
				`rsct.compat.clients.sp 2.4.5.0`
				`rsct.core.auditrm 2.4.5.0`
				`rsct.core.fsrm 2.4.5.0`
				`rsct.core.gui 2.4.5.0`
				`rsct.core.hostrm 2.4.5.0`
				`rsct.core.lprm 2.4.5.0`
				`rsct.core.sec 2.4.5.0`
				`rsct.core.sensorrm 2.4.5.0`
				`rsct.core.sr 2.4.5.0`

Parallel
ESSL
	320	5765F8400		`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`

ESSL	420	5765F8200	[PK02186]	`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`

Recommended Installation Sequence

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

Notes to IBM SSR's and System Administrators

In order to co-ordinate 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:

- [System p 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 [System p 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 System p Maintenance Guide](#)
- [Hardware Management Console for System p Installation and Operations Guide](#)
- [System p 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 as appropriate :

- SF230: Version 4.5 machine code updates
- SF235: Version 5.1 machine code updates
- SF240: Version 5.2 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](#).

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

Scroll to "Supported HMC products" and select as appropriate :

- SF230: Version 4.5 machine code updates
- SF235: Version 5.1 machine code updates
- SF240: Version 5.2 machine code updates

Follow the information and HMC installation instructions.

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: 5
Release: 2.0
HMC Build level 20060210.1

Level Check from HMC GUI
Menu "Help" > "About Hardware Management Console" shows:

230:
Version: 4
Release: 5.0
HMC Build level 20051010.1
MH00454: Required Maintenance Fix for V4R5.0 (10-10-2005)

235:
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:
Version= Version: 5
Release: 2.0
HMC Build level 20060210.1
MH00586: Required fixes for HMC V5R2.0 (02-14-2006)
MH00610: Fix Dump collection issues (03-09-2006)

Installation Notes
Whenever you update HMC software, you must subsequently either reboot the HMC or rerun the addpeer
and link commands to reconfigure the updated HMCs for CRHS.

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.

If you upgrade the HMC version, re-run the addpeer command on the management server, and reassign
frames to the HMC to avoid losing the the HMC peer domain setup.

Have your IBM SSR download the 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 System p Installation and Operations Guide](#).

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 imbedded kernel will not run well when this option is enabled. This setting 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

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

This may take up to 10 minutes to complete.

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

HMC Properties

Additional information on the HMC properties is also available in Chapter 6, "Installation" of the System p High Performance Switch (HPS) Planning, Installation and Service Guide for IBM eServer p5 servers

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

Setting 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 System p High Performance Switch (HPS) Planning, Installation and Service Guide for IBM eServer p5 servers (SA38-0646-00) for more details.

Enable CRHS Port
Enable the CRHS port

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

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.
- Reboot HMC after the configuration

Set Cluster 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 System p 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 System p 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 frames must be in the Standby/Standby state on the "Frame Management" panel prior to the power subsystem microcode install/upgrade.

Documentation Reference:

- [System p 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 [System p 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:

- [System p 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.
- 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: [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 Assignment

Create LPAR definition and assign adapters.

For details refer to "HMC code load" in the "Software Installation" section in Chapter 6, "Installation", System p HPS Planning, Installation and Service Guide.

Install Cluster System Manager (CSM)
CSM software is part of the base AIX installation. For the recommended AIX service levels see Compoment Information

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
- System p 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 System p High Performance Switch (HPS) Planning, Installation and Service Guide for IBM eServer p5 servers (SA38-0646-00).**

b. Apply RSCT and CSM Updates to CSM MS

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

RSCT PTF Download Information - refer to Component Update/Download Information for RSCT

CSM PTF Download Information - refer to Component Update/Download Information for CSM

Level Check:

- ❍ lslpp -L csm.client csm.core csm.diagnostics csm.dsh csm.gui.dcem csm.gui.websm csm.server csm.hams
- ❍ Refer to CSM Level Check
- ❍ lslpp -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
- ❍ <>refer to AIX52 RSCT level check
- ❍ refer to AIX53 RSCT level check

Define the nodes/LPARs in the cluster.
- 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 System p 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 [System p 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 should point 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)]()
- [System p 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]()
- Also contains details on installing the AIX operating system on the LPARs.

AIX5L Maintenance Package Download Information:

To download:

- Go to [http://www-03.ibm.com/servers/eserver/support/unixservers/aixfixes.html]()
- Select Maintenance Packages and choose AIX52 or AIX53
- Specify your "Current level"
- Specify "Desired level" and click "GO"
- Follow instructions on this page.

Also check AIX5L Critical Fixes.

Level Check:

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

- oslevel -r
- lslpp -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 Update images to all LPARs
Note that after 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:

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

Expected Levels are give 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 LoadLeveler 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.

Setup 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..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

Setup 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.

EX: 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.

EX: /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

Notes

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.

Checking the Large Page Option
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(bytes of real mem) = T(bytes of reaL mem allocated to TLP) + S(bytes of real mem allocated to Small Pages)
- Small Page memory is a function of S(bytes of real mem allocated to Small Pages) = N(bytes of real mem) - T(bytes of reaL mem allocated to TLP)
- Large Page memory is a function of T(bytes of reaL mem allocated to TLP) = N(bytes of real mem) - S(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 Adapter Switch Network Interfaces
Refer to [System p (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)
  a. Disable Dynamic Host Configuration Protocol (DHCP) on the HMC and reboot the HMC.
  b. 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](#)
  c. 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. 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 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 System p 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:

- lslpp -L csm.client csm.core csm.diagnostics csm.dsh csm.gui.dcem csm.gui.websm csm.server csm.hams
    - AIX52 / AIX53 CSM Level Check
- lslpp -L rsct.core.sensorrm rsct.core.utils rsct.exp.cimrm rsct.basic.rte rsct.core.lprm rsct.core.rmc rsct.core.sec rsct.basic.hacmp rsct.core.auditrm
    - AIX52 RSCT Level Check
    - AIX53 RSCT Level Check

Notes:

a. Please see the following CSM Webpage for packages and known issues: http://www14.software.ibm.com/webapp/set2/sas/f/csm/download/home.html
b. The CSM Management Server should have the right levels of RSCT and CSM PTFs after 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 with Step 13a.
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 [System p High Performance Switch (HPS) Planning, Installation and Service Guide for IBM eServer p5 servers.](#)

Download Information:

❍ Download the HPSNM Service Pack package in the HPSNM updates and fixes section of:
❍ [HMC POWER5 page](#).
❍ Download csm.hpsnm AIX PTF from
❍ 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_Serv_Pack_2.zip 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 System p 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:

❍ System p 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 lslpp -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.

Restrictions/Known Problems/Workarounds

This System p 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 System p HPS POWER5 Service Pack and try to enable the redundant FSP feature with invalid Klingon configurations may result in ambiguous to troublesome errors.
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. 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, 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:
Whenupdating 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
- 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:

D103BOFF Predictive Error

Problem Description:

An informational error was incorrectly categorized as a predictive failure

Customer Impact:

Unnecessary call for service action

Workaround:

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

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
   - ❍ Vheck 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 is to fix issue in a subsequent Service Pack.

HPS Service Pack Abstract - By Component
- [HMC](#)
- [SNM](#)
- [GFW](#)
- [VSD](#)
- [LAPI 233](#)
- [LAPI 242](#)
- [HPS 113](#)
- [HPS 120](#)
- [PPE 411](#)
- [PPE 422](#)
- [LoadL 320](#)
- [LoadL 331](#)
- [GPFS 210](#)
- [GPFS 221](#)

- [GPFS 230](#)
- [CSM 15](#)
- [RSCT 238](#)
- [RSCT 244](#)

HMC
Please refer to the [HMC website](#) for specific details.
SNM
N/A
GFW
N/A
VSD
4.1.0.14

On purple at LLNL we're seeing that the trace file for IBM.vsd
VSD.SNAP GETS ERROR WHEN REMOTE FS NAMED /TMP IS MOUNTED
VSD.SNAP ERROR IN THE SUBROUTINE GET_CSS_STUFF
NODE PANIC IN VSD DURING FAILOVER IN VSDD:SUSVSD.
VSD 2-node recovery is failing
4.1.0.15

panic:SndRmt vsdkt threadentry
VSD system panic in RcvRmt
LAPI
2.4.2.5

INCORPORATE NEW LAPI ENV VARIABLES IN NEXT SERVICE
2.4.2.4

LAPI assert running dual protocol mpi/lapi in striped mode
2.3.3.5

LAPI interrupts disabled when MP_MSG_API=MPI,LAPI
HPS/SNI
1.2.0.3

64K page support
Preemption issues
RCXT_READ_IOCTL
med page test cases

1.2.0.2

SNIDD PANIC IN XMEMDMA64()
CRASH IN MLTDD
ASC purple: zdebugdrv fileset needs a post_u
1.1.3.7

SNIDD PANIC IN XMEMDMA64()
CRASH IN MLTDD
1.1.3.8

64K page support
Preemption issues
RCXT_READ_IOCTL
Parallel Environment
4.2.2.3

MP_PRINTENV PRINTS MISLEADING INFO ABOUT MP_INSTANCES
FAILURE TO CREATE LIGHT WEIGHT CORE FILES.
Update AIX Trace hooks in MPCI
INCORPORATE NEW LAPI ENV VARIABLES IN NEXT SERVICE
timing issues in rset attachment
4.2.2.2

SAMPLES SCRIPT /USR/LPP/PPE.POE/SAMPLES/AFS/BUILDAFS FAILS TO
POE GENERATES NULL CHARACTERS ON MP_PRIORITY=YES - FIX REQUI
POE FAILS WHEN TOTALVIEW FAST CONDITIONAL BREAKPOINT USED
THE PARTITION MANAGER DAEMON (PMDV4) TERMINATED WITH A CORE DU
ASCI: MP_BUFFER_MEM second argument being ignored
4.1.1

POE CANNOT INHERIT LARGE LIMITS
Loadleveler
3.2.0.18

IN THE LOADL JCF FILE THE DEPENDENCY STATEMENT IS ONLY 1024 CH
WHEN USING API SCHEDULER AND USING LLMODIFY CMD TO CHANGE THE
WHEN A MACHINE IS REMOVED FROM AN LL CLUSTER BUT IS STILL RUIN

3.2.0.19

Problem with preemption in presence of adapter failures/errors.
INVALID OPTION COMBINATIONS FOR LLQ COMMAND SHOULD FAIL WITH E
LOADL_CONFIG VARIABLE NOT WORKING PROPERLY.
JOB NOT REJECTED IF INITAIL DIRECTORY IS MISSING WHEN RUNNING
LOADL_CONFIG VARIABLE NOT WORKING PROPERLY.
LYN build breaks in building on the changes in the track 133815

3.3.2.1

INVALID OPTION COMBINATIONS FOR LLQ COMMAND SHOULD FAIL WITH E
PREEMPT TIMES NOT PRESERVED IF NEGOTIATOR DOWN/UP-WALLCLOCK HI
WHEN A JOB IS MOVED FROM ONE LOADLEVELER CLUSTER TO ANOTHER VI
BGL: LoadL_Neg coredumps with V1R2M1_020_2006-060110/ppc build
Memory error in query api for CLASSES
WHEN SPEED SET TO LESS THEN 1 IN ADMIN FILE THE JOB CPU IS CAL
llq -s top dog message
JOB CONTROL FILES ARE NOT CHANGED AS SPECIFIED BY THE REMOTE C
RHEL 4 X86: file /tmp/startd_unix_dgram_socket is not deleted.
LL signals to mpirun not interpreted correctly
THE PRESENCE OF FILE LOADL_CM WILL CAUSE PROBLEMS LLSTATUS AND
LL jobs do not start if ownr field exists & PART STATE=FREE
Need to update loadleveler.pm for IBM globus
Need to update GT4 document.
JOB STEP ID MACHINE NAME PART MISSING FROM LLSUMMARY OUTPUT.
Neg. crash during user constraints sys initiated premption run
LL keywords processing nonnumeric characters w/o errors
Doing ctrl-C on down node during "llctl -g start" exits command
2EB:LoadL_starter coredumps due to corrupt stack
2EB :llstatus -a shows "NOT READY" with dynamic adapter setup
Linux: llextRPD missing from PPC/Red Hat, Opteron builds
2EB: ifconfig ib0 down, llq -s does not report reason
2EB: LoadL_startd crashes processing incorrect info from RSCT
Message Catalog Updates
LL Mercury PTF1 needs to show correct level 3.3.2.1
LL 3.3.2 shows no adapters w/rsct at 2.3.5
starter crashed during checkpoint command
defect created to remove changes from 133937

GPFS

2.3.0

KERNEL HANG IN KSFSREMOTERETRY AND KXSENDFLOCK

2.2.1

N/A

2.1.0

DCR FOR CONFIGURING ML0 TO DELETE ML ROUTES
CSS PING LATENCY WHEN IPSEC IS ENABLED
DSI CRASH DURING BOOT WITH CRUISER/CORSAIR A
CS_SW_GET_SVCREQ_ER NODE OFF SWITCH
/ETC/PASSWD AND GROUPS PERMISSIONS NOT SET B
LAPI_INTERRUPT_OFF
DCR REQUEST FOR CSS IP DRIVER TO RUN AT INTC
CSM
1.5.1.1

DCmonitorinstall():current_status and status_history need added
rfwflash does not work with JS21 blades
dsh fails for devices defined with ip address
attribute validity checking needed for clusterinfo and nodeinfo
DC.log should not get created in CSM environment
copycds need to prompt user if no CD or ISO mounted
mgmtsvr fails when the managed node has csm.server installed
A typo "x64_64" in csmsetupks manpage
Confusing getadapters messages from csmsetupyast
Need to add Click-to-accept to FSP direct attach
RHEL: installing an i386 node from an x86_64 MS failed
DSH NO LONGER PASSES -N TO RSH WHICH BREAKS PIPES
'systemid [-c] -s' does not work
AIX/IVM:SLES9SP3 installation failed when installing packages
IBGA7ST:csmstat shows PowerStatus unknown
THE CSMSETUPNIM COMMAND CHECKS FOR HOSTNAMES ON NODE GROUPS TH
getadapters auto stanza file cannot be read
Scripts in csm.ivm.server and csm.ivm.client are not running
setbootdisk does not work on JS20 with --noreboot
HMC not disconnecting stream server in all required cases
HAMS started,setupyast failed while using UUIDmethod for instal
support direct cross subnet installs
Update CSM 1.5.1.1 README
getatapters does not set machine_type=install
Update VRMF for CSM 1.5.1.1
mgmtsvr fails when the managed node has csm.server installed
installms/copycds/copycsmpkgs does not work with the real CDs
AIX: Rejecting csm.server 1.5.1.1 breaks root part
RSCT

2.4.5.1

WEBSM REPORTING WRONG NODE DOWN FOR SWITCHRESPONDS
LX: Local disk isn't harvested when machine has shared disks
0HK: Name field should better not contain /
CT: 22 out of 112 nodes ctrmc become inoperative after update
Name of AgFS should change to resource id when fsid/label blank
2EB: isNodeStopped is not reset when joining
LX: reset disk results in agfs opstate change to 4
0A0: Code Review Comments - Halt Callback
0A0:5 seconds delay of notification for hags
ConfigRM fails occationally on startup
OpState can hang in the Pending Online state
Intermittent: processes dealing with IBM.Partition hang
Failure to disable cthagsglsm when no HPC adapters exist
0HK AX: Core.0 creates when online AgFS
Source and destination overlap in memcpy()
0HK: re-wording of some messages in IBM.StorageRM.msg
RSCT when adapter is bad loops resending down notice to domain
0HK AX: Reset does not take VG offline and leave OpState=4
OpState stuck online after daemon bounce is ONLINE
Allow including libsysfs.h without changing it
KMUX: A memory leak bug
IBGA7ST: /sbin/srcmstr Segmentation fault
LX: RSCT Debussy man page updates
IBGA7ST:lsrcrc IBM.NetworkInterface OpState show more interface
THL update for address swap between two interfaces fails
CPU USAGE INCREASES WITH RSCT 2.4.1.4
increment release # for ppc linux
Set VRMF for Debussy/Debussy52 PTF 1
Harvest appears to leak memory
Different trace category causes the improper trace output
Improper trace category - continue 134938
0A0: Disk registration causes crash with 64-bit kernel
hats sends duplicate incarnation numbers to hags
2EB: rsct info on ib interfaces missing on 24 nodes

2.3.9.1

WEBSM REPORTING WRONG NODE DOWN FOR SWITCHRESPONDS
LX: Local disk isn't harvested when machine has shared disks
0HK: Name field should better not contain /
CT: 22 out of 112 nodes ctrmc become inoperative after update
Name of AgFS should change to resource id when fsid/label blank
2EB: isNodeStopped is not reset when joining
LX: reset disk results in agfs opstate change to 4
0A0: Code Review Comments - Halt Callback
0A0:5 seconds delay of notification for hags
ConfigRM fails occationally on startup
OpState can hang in the Pending Online state
Intermittent: processes dealing with IBM.Partition hang
Failure to disable cthagsglsm when no HPC adapters exist
0HK AX: Core.0 creates when online AgFS
Source and destination overlap in memcpy()
0HK: re-wording of some messages in IBM.StorageRM.msg
RSCT when adapter is bad loops resending down notice to domain
0HK AX: Reset does not take VG offline and leave OpState=4
OpState stuck online after daemon bounce is ONLINE
Allow including libsysfs.h without changing it
KMUX: A memory leak bug
IBGA7ST:lsrcrc IBM.NetworkInterface OpState show more interface
THL update for address swap between two interfaces fails
CPU USAGE INCREASES WITH RSCT 2.4.1.4
Set VRMF for Debussy/Debussy52 PTF 1
Harvest appears to leak memory
Different trace category causes the improper trace output
Improper trace category - continue 134938
0A0: Disk registration causes crash with 64-bit kernel
hats sends duplicate incarnation numbers to hags
2EB: rsct info on ib interfaces missing on 24 nodes