

- [Cluster support](#)
- [HPS](#)

IBM HPS POWER 4 Service Pack 22 Addendum 1 FLASH/Readme

Introduction

Anchor link navigation

- **[Introduction](#)**
- [Supporting documentation](#)
- [Software updates](#)
- [Component updates/download information](#)

- [Installation sequence](#)
- [Installation guidelines](#)
- [Known problems and workarounds](#)
- [Fix list and readme archive](#)

Introduction

12/19/2007

IBM pSeries High Performance Switch (HPS) Service Pack 22 Addendum 1 supports new AIX and Cluster LP release and service levels, as well a new HPSNM service level for the HMC. There were no changes to the GFW, Power Code or HMC levels in this addendum.

This Service Pack details updates to these components:

- [Hardware Management Console \(HMC\)](#) 3.3.7 - Build level 20060207.1
- [Global Firmware \(GFW\)](#) 3H070425 / 3J070425
- [Power Subsystem Microcode \(PT Code\)](#) 26A9
- [Switch Network Manager \(SNM\)](#) 1.3.9.0
- [AIX 5L Version 5.3](#) 5.3.7.1 (5300-07-01-0748)
- [AIX 5L Version 5.2](#) 5.2.0.109 (5200-10-03-0744)

[AIX High Performance Computing \(HPC\) and Clusters LPs](#) as listed in the [Detailed LP Level Check](#)

- [VSD](#)
- [LAPI](#)
- [HPS/SNI Devices](#)
- [PPE](#)
- [LoadLeveler](#)
- [GPFS](#)
- [CSM](#)
- [RSCT](#)

Note: These recommended levels should not be considered "minimum", unless specifically stated as such in this document. Also, always refer to the individual component support/download sites for the latest service bulletins, security alerts, prerequisites, known problems, and restrictions relating to that product.

This Service Pack also contains general guidelines for upgrading the components listed in the **Component updates/download information** section.

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 the **Supporting Documentation** section. The Code Levels listed in the **Component updates/download information** section reflect the levels available at the time of this Service Pack 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 **Installation Sequence** section is the standard sequence of installation. Non-standard sequences or undocumented code levels may cause unforeseen problems. In this event please contact your Customer Service Representative.

Because '/var' is a system data repository, system administrators should check periodically to maintain /var such that there is at least 30 Mb free (use the **df -k** command). If it is more than 75% full, look for the directories that contain the most data (use the **du /var | sort -n** command).

Health check/planning help

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

- [Cluster support](#)
- [HPS](#)

IBM HPS POWER 4 Service Pack 22 Addendum 1 FLASH/Readme

Supporting Documentation

Anchor link navigation

- [Introduction](#)
- **[Supporting documentation](#)**
- [Software updates](#)
- [Component updates/download information](#)

- [Installation sequence](#)
- [Installation guidelines](#)
- [Known problems and workarounds](#)
- [Fix list and readme archive](#)

Supporting documentation

Refer for to the following documentation for additional information to supplement the HPS Readme.

- [Hardware Management Console](#)
- [Hardware Management Console for pSeries Installation and Operations Guide](#)
- [Hardware Management Console for pSeries Maintenance Guide](#)
(SA38-0603-05)
- [pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide](#)
(GA22-7951-01) (HPS Guide)
- [IBM eServer pSeries 690](#)
- [Switch Network Interface for eServer pSeries High Performance Switch Guide and Reference](#)
(SC23-4869-01)
- [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 for AIX 5L Installation Version 4 Release 1.1](#)
- [IBM LoadLeveler for AIX 5L and Linux Using and Administering Version 3 Release2](#)
- [LoadLeveler 3.2 documentation updates](#)
- [pSeries and AIX Information Center](#)
- [AIX 5L Version 5.3 documentation](#)
- [System management guides](#)
- [Installation guides](#)
- [HPS POWER4 Readme Archive](#)

Health check/planning help

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

- [Cluster support](#)
- [HPS](#)

IBM HPS POWER 4 Service Pack 22 Addendum 1 FLASH/Readme

Software Updates

Anchor link navigation

- [Introduction](#)
- [Supporting documentation](#)
- **[Software updates](#)**
- [Component updates/download information](#)

- [Installation sequence](#)
- [Installation guidelines](#)
- [Known problems and workarounds](#)
- [Fix list and readme archive](#)

Software updates

The following software updates are included:

- [New SNM/FNM features](#)
- [MPI and LAPI performance enhancements](#)
- [Support for Node Switch Board \(NSB\) and Intermediate Switch Board \(ISB\) failure](#)
- [Application Striping and Fail-Over of packets across multiple interfaces for fault resilience](#)
- [Application Checkpoint/Restart and preemption](#)
- [The Web-based System Manager Remote Client \(WebSM PC CLIENT may need to be reinstalled](#)
- [RDMA \(remote direct-memory access \) transport on HPS](#)

New SNM features

This section discusses the new features and functions on the Switch Network Management Panel introduced in HMC 1.3.1.0 Service Pack 9.

For details of operation please see Chapter 4. System management components and Appendix E. Switch Network Manager (SNM) in [pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide \(GA22-7951-02\) \(HPS Guide\)](#)

Features

HMC Fail-over:

- FNM/SNM daemon Enable SNM Software in Normal Mode runs on up to four HMCs in the cluster.
- Verification mode: Enable SNM Software for Switch Network Verification used to check out the cluster.

New features:

- Select Logical Topology
- Display Cluster Components
- Enable SNM Software for Switch Network Verification

The "Enable SNM Software for Normal Mode" has been renamed. It used to be "Enable SNM Software"

Select Logical Topology introduces a requirement to set the logical topology of your system before enabling SNM software. If the logical topology has not been set, neither of the tasks to Enable SNM Software will be selectable. This is required. See note 1.

Note: If you are using the SNM Fail-over support introduced in the Software Update release, the logical topology must be set on each HMC where you enable SNM software.

Display Cluster Components shows the frames and cages of the cluster components with which SNM can communicate. This task can only be used if SNM software is *not* enabled. The lsswcomp command provides the same function as the "Display Cluster Components" task.

Enable SNM Software for Switch Network Verification task is used to check out the cluster by NOT removing bad links or routes. The main purpose for this feature is to discover cable mis-wires and defective hardware. No Routing is modified or loaded. It is NOT intended to be used in a working environment.

Enable SNM Software for Normal Mode. Select this feature to run SNM in a normal, working environment. This is required. See note 1.

Notes:

1. If you are upgrading from Service Pack 8 or less, these steps are *required* for an initial installation of FNM/SNM:

First Select Logical Topology, and then select Enable SNM Software for Normal Mode. See "Enable SNM Software" in the HPS Guide.

2. The systems topology selected needs the number of NSB's and ISB's calculated on a per plane basis.

Example: For a configuration with two NSB's and two planes, there is one NSB per plane with 16 endpoints per plane. This equates to 1 Network, 2 Planes, and 1NSB_OISB_16EP.

[Back to top](#)

LAPI Performance Enhancements

LAPI and MPI host communication stacks have been tuned, based on experience, for lower latency. These changes complement the communication performance improvements released in part 1. They are being released as one package with the LAPI and MPI striping function.

Note: Beginning with the Service Pack 9:

- A single MPI user-space job which wants to use HPS adapters in both of two switch planes must set `MP_EUIDEVICE=sn_all`(or `csss`). In previous releases, a single MPI/us job with multiple tasks per node could use adapters from two planes by setting `MP_EUIDEVICE=sn_single` (or `css0`).
- A job which sets `MP_EUIDEVICE=sn_single` will only be able to use half of the HPS adapters on the nodes where it runs.
- A job must set `MP_EUIDEVICE=sn_all` in order to use all the HPS adapters in the nodes it is using.

Support for Node Switch Board (NSB) and Intermediate Switch Board (ISB)

The SNM daemon has been enhanced to ensure that the failure of an NSB and an ISB will be handled appropriately. The main enhancement is that the Switch Network Manager daemon gathers and downloads multiple path table updates into a single transaction instead of one update per transaction.

An entire switch board failure or recovery will generate multiple of path entries changes. Therefore handling these updates in groups involves many fewer transactions. These can be processed in a more

timely fashion and place a smaller load on SNM and the service network.

Application striping and fail-over of packets across multiple interfaces for fault resilience

Striping provides a method for a single task of a parallel application to utilize multiple SNI links. This allows a single link to fail without the application using the link being terminated because there is an alternate path to all other tasks of the application. The striping method implemented is designed to provide resilience to switch or interface/link failure and is not designed to provide an increase in the aggregate network bandwidth as seen by tasks.

This link failure resilience function requires support in LoadLeveler, Parallel Environment and the host communication stack for MPI and LAPI. When an interface/link recovers, the striping function will ensure that the link is used again. Another goal of this striping design is to evenly distribute traffic over multiple switch networks to maintain a balance in the communications load.

Known issues and concerns

Known issues and concerns include the following:

- In non-striping mode you only get half the links on a p655 system using a 2 plane configuration when setting `MP_EUDEVICE=sn_single`. You must use `MP_EUDEVICE=csss` or `sn_all` to get all links.
- Performance degradation for single task per link of 3% for unidirectional and 5% for exchange bandwidth at large messages. However, multiple tasks per link get full link bandwidth.
- Striping currently limited to two links per task. Multiple tasks can use all links as long as there are more tasks than half the links available.

Note: When you are using RDMA, there is no large message performance penalty for striping. In fact striping greatly improves large message RDMA bandwidth. The above performance degradation occurs only in packet mode.

[Back to top](#)

Application Checkpoint/Restart and Preemption

Checkpoint/restart provides a way to stop and resume applications at some later time. It is a very useful tool in managing the clusters workload since applications can be check-pointed to allow other applications to run or the system made inactive for a maintenance window.

Preemption is also provided so a running application can be suspended to allow another higher priority application to execute right away.

The Web-based System Manager Remote Client (WebSM PC CLIENT)

If the Web SM is updated, it is recommended that the existing code be reinstalled on your remote server or PC to ensure that the WebSM is compatible with this HMC Build.

To install the PC Client software on your remote server or PC go to: http://<hmc-hostname>/remote_client.html.

The two versions are the "legacy" WebSM client - "InstallShield" and "The future of WebSM clients." - Webstart. Both versions can exist simultaneously on your client workstation.

Install Shield

If there is a currently installed install shield version of WebSM, it *must* be **uninstalled** before the new version is installed.

To uninstall, follow these steps:

1. From the Start Menu, select Control Panel -> Add/Remove Programs.
2. Select Web-based System Manager Remote Client, and then click "Change/Remove".
3. Follow the on screen instructions.

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

To install, follow these steps:

1. Go to http://<hmc-hostname>/remote_client.html.
2. Select "Install Shield".
3. Select the Operating System where the program will be running.
4. Select "Open" to immediately install on the machine on which the browser is running.
5. Selecting "Save" allows you to store the install program for installation at a later time.

Note: The Install Shield package is a 100 Mb file, so either process will require about 15 minutes. Once saved, the self-extracting file can be used to upgrade multiple PCs.

Webstart

The webstart installation is a two part process:

1. Installing the 1.4.2 JVM that contains the webstart code.
2. Invoking the launch code and gets the initial classes downloaded from the HMC.

When the client is launched, it checks for new jar files on the HMC and downloads them, so subsequent updates of the HMC server code do not require the updates of the webstart client code.

Having multiple jvms installed on the client workstation may result in the default jvm path not pointing to the IBM 1.4.2 jvm. Since webstart does not work well with the Sun 1.4.2 jvm, you may need to clean up multiple jvm installations on your workstation.

The Java Web Start version requires multiple logins on launching. You have to log in to the original HMC whenever you try to connect even if you are just downloading the webstart **.jsp** files. Downloading the **.jsp** files from a regular AIX machine requires one less login.

Multiple logins are not an issue in the non-webstart websm because you are not connecting to the webserver to download the .jsp files. Trying to 'cancel out' results in multiple (up to 16) presentations of the login verification box.

1. Install Java Web Start:
 - o Web Start for Linux- Install Java Web Start on a Linux platform. Once installed, return to this page to download the Remote Client.
 - o Java Web Start for Windows- Install Java Web Start on a Windows platform. Once installed, return to this page to download the Remote Client.
2. Download the Web-based System Manager Remote Client for Java Web Start on Linux and Windows systems.

[Back to top](#)

RDMA transport on HPS

The IBM eServer pSeries High Performance Switch (HPS) with Remote Direct Memory Access (RDMA) enabled delivers 1.8GB/s of unidirectional bandwidth across a single link and over 3.5GB/s over two links to a single MPI task for some applications. Bidirectional bandwidth measures almost 3.0GB/s over a single link and over 5.9GB/s for a single MPI task striping over two links.

These measurements were achieved on two nodes containing 1.9GHz, Power 4+ based processors with 32 CPUs and two HPS adapters, and links on each of the two nodes. Large pages and cache-line aligned buffers were used for the MPI measurements. These results indicate nearly linear scaling of large message striping performance. A single link is able to deliver over 1.65GB/s of bandwidth over TCP/IP.

Advantages of RDMA

RDMA provides the following advantages:

- Decouples the CPU from the movement of data allowing for better overlap of computation and communication.
- Reduces the computational load of the CPU by off-loading segmentation and reassembly of messages to the network adapter, thereby reducing the number of packet arrival interrupts.
- Reduces the stress on the memory subsystems by reducing the number of bus crossings - one I/O bus crossing versus the traditional two memory bus and one I/O bus crossing when not using RDMA.
- Enables the protocols to efficiently stripe a message (or different messages) from a single task across multiple network interfaces to exploit the available communication bandwidth in parallel without engaging multiple CPUs.
- Provides improved raw transport performance. In cases where the transport bottleneck is the copy rate (memory bandwidth), RDMA helps eliminate that bottleneck.

How application can exploit RDMA

To exploit the overlap potential of RDMA, applications should make use of non-blocking calls and schedule communication as early as possible to extract the full overlap potential. For MPI/LAPI applications the RDMA transport kicks in for large messages (to justify the setup costs of RDMA).

Applications that reuse the same communication buffers will benefit further from RDMA since the setup cost is amortized over larger number of transfers. For applications that send large messages asynchronously, enabling striping should help the applications realize superior transport bandwidth.

RDMA transport is **disabled** by default in this code release, to insure that customers running production workloads that do not need RDMA operate without change.

Note:

When RDMA is enabled by changing the device attribute **rdma_xlat_limit** to a value greater than 0, it is possible for an RDMA enabled job to surpass the AIX pin limit which may lead to a hung system. If this pin limit is exceeded slowly enough the parallel application will catch a SIGDANGER signal and kill the job. However, if this pin limit is exceeded too rapidly, AIX may resort to killing processes, including systems processes, which may cause system to appear to be hung. This situation may be prevented by setting the **rdma_xlat_limit** to a number which is smaller than 4096 times the sum of the number of large and small pages on the system.

RDMA usage

By default, the RDMA capability of the SNI adapters is turned off. Use the `chdev` (or optionally `/usr/sni/aix52/chgsni`) command to turn on RDMA by setting `rdma_xlat_limit` to a value greater than `0x0`.

The `rdma_xlat_limit` attribute controls amount of memory allowed for RDMA, per LPAR. The default value of the `rdma_xlat_limit` attribute is `0x0`, effectively turning RDMA off.

When using RDMA it is recommended that this attribute value be set to value equal to 75% of small page memory. This will allow all of the large pages on the system and most of the small pages to be mapped to the sni adapters.

Note:

This odm attr (`rdma_xlat_limit`) is changed for ALL sni's in an LPAR with just the single command. For the change to take effect, the LPAR must be rebooted. The RDMA option MUST be turned ON OR OFF across the entire cluster. If it is not, SNI adapters will not communicate due to version mismatch failures.

When RDMA is turned on, the IP protocol automatically uses the RDMA capability. For user space jobs, an additional LoadLeveler keyword (`bulkxfer=yes`) must be set in the LoadLeveler job control file to indicate that this job is requesting RDMA. This allows users to run both RDMA and non-RDMA jobs on the system.

For the most current details, refer to [LoadLeveler 3.2 documentation updates; Addition of support for Bulk Data Transfer](#).

For additional information, refer to [LoadLeveler Using and Administration Guide](#).

To enable RDMA:

```
/usr/sbin/chdev -l sniX -a rdma_xlat_limit=XLAT_LIMIT
```

where `XLAT_LIMIT` is 75% or less of the small page memory.

To disable RDMA:

```
/usr/sbin/chdev -l sniX -a rdma_xlat_limit=0x0000000000000000
```

To check the RDMA status: -- run `lsattr -El sni0 ; lsattr -El sni1 ;`

Example:

```
for i in `lsdev | grep sni.*Avail | awk '{print $1}'`; do lsattr -El $i|grep rdma; done
```

`chdev Usage: chdev -l Name [-a Attribute=Value]...[-p ParentName][-P|-T]`

```
chdev -l Name -a rdma_xlat_limit=size
```

The 'True' seen in the lsattr output is whether or not the value is user modifiable. The rdma_xlat_limit attribute is one of 5 that are not valid on just 1 sni. The change method updates the odm for all snis.

This attribute is modified for all existing devices. Any user supplied -l option is ignored.

Attribute values for rdma_xlat_limit :

Default value: 0x0 (Implies RDMA is turned off)

MINimum value: 0x0

MAXimum value: 0xFFFFFFFFFFFFFFFF (Maximum amount of memory allowed for RDMA (in bytes))

Recommended value: 75% or less of small page memory.

[Back to top](#)

Regarding the use of memory by RDMA and TLP

Real memory is divided into two categories - Small Pages and Large pages. It is the user's responsibility to achieve an optimal balance between the the two categories based on the expected and/or experienced needs of both SNI adapters memory requirements expressed in TLP and applications use of Small Pages as expressed in RDMA. TLP can allocate up to 75% of real memory. RDMA can pin and map up to 75% of small page application memory.

Total Real Memory is a function of $N(\text{bytes of real mem}) = T(\text{bytes of real mem allocated to TLP}) + S(\text{bytes of real mem allocated to Small Pages})$.

Small Page memory is a function of $S(\text{bytes of real mem allocated to Small Pages}) = N(\text{bytes of real mem}) - T(\text{bytes of real mem allocated to TLP})$.

Large Page memory is a function of $T(\text{bytes of real mem allocated to TLP}) = N(\text{bytes of real mem}) - S(\text{bytes of real mem allocated to Small Pages})$.

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

lsattr -E -l sys0 -a realmem returns the number of kbytes real memory, call this number A.

vmo -o 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.

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 it executing. After it exits the memory is unpinned and freed.

See also [Regarding the use of memory by TLP](#).

[Back to top](#)

Health check/planning help

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

- [Cluster support](#)
- [HPS](#)

IBM HPS POWER 4 Service Pack 22 Addendum 1 FLASH/Readme

Component Updates and Download Information

Anchor link navigation

- [Introduction](#)
- [Supporting documentation](#)
- [Software updates](#)
- [**Component updates/download information**](#)

- [Installation sequence](#)
- [Installation guidelines](#)
- [Known problems and workarounds](#)
- [Fix list and readme archive](#)

Component updates

This section lists the service packs for the various components and provides information or links for downloading these service packs. It also details which releases of AIX LPs (components) are compatible with which AIX release and provides detailed level checks for the AIX LPs (components).

Anchor link navigation

- [Service pack and download information](#)

- [AIX Licensed Products](#)

Service pack and download information

The following table lists service pack information for each component and provides links for downloading or ordering the service packs. In addition, the Component name link displays installation

information for the component.

Component	Service pack updates	Download sites
HMC	<p>The HMC code level for this Service Pack is HMC V3R3.7 Choose HMC_Update_V3R3.7.zip APAR IY82939 PTF U807279</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>Note: If you plan to use the removeable media, please read the Important Notice regarding the use of the Format Removable Media task: HMC POWER4 updates.</p>	<p>HMC corrective service OR HMC POWER4 updates</p>
SNM/FNM	<p>The SNM code level for this Service Pack is SNM Version: 1.3.9.0-1. STATUS RPM Version: 1.3.9.0 SNM Version: IBMhsc.SNM-1.3.9.0-1.i386</p>	<p>HMC corrective service OR HMC POWER4 updates</p>
GFW	<p>The Firmware level for this Service Pack is ROM Level (alterable)..... 3H070425 (P690) ROM Level (alterable)..... 3J070425 (P655)</p> <p>This code is machine specific. Use option 1, Download microcode by machine type and model 1) Download microcode by machine type and model</p> <p>For p690 models: Select "7040-681" and then choose "Version 3H070425"</p> <p>For p655 models: Select "7039-651" and then choose "Version 3J070425"</p> <p>Note: GFW is available in IBM CORE 3-4 days earlier than the above mentioned website. Please contact your IBM CE for the GFW in IBM CORE if not available on website.</p> <p>Refer to this website for detailed download and unpacking procedures:</p>	<p>Microcode downloads</p>

[Microcode downloads](#)

[GPFS](#)

For the latest information on migration, coexistence, and compatibility, see the [GPFS FAQ](#).

[GPFS Support and downloads](#)

[Power Subsystem Microcode](#)

The Power Subsystem Microcode level for this Service Pack is 26a9.

[Microcode downloads](#)

Go to "Power Subsystem for 7039-651 (p655) and servers containing the 7045-SW4 (High Performance Switch)."

Code is the same for both p690 and p655:
ptcode-1.78.26a9-1.i386.rpm

Note: Power Subsystem Microcode is available in IBM CORE up to one week earlier than the above mentioned website. Please contact your IBM CE for the Power Subsystem Microcode in IBM CORE if not available on website.

[AIX 5.3](#)

The AIX 5.3 level for this Service Pack is 5.3.7.1

[System p support fixes](#)

AIX53: 5300-07-01-0748

[AIX 5.2](#)

The AIX 5.2 level for this Service Pack is 5.2.0.109

[System p support fixes](#)

AIX52: 5200-10-03-0744

CSM 1.4 and above requires RPM update openCIMOM 0.8 (5.2)

CSM openCIMOM

If you are APPLYING software for csm. server 1.4.1.1 or higher, please install the following images from the AIX Installation Media Volume 2:

[openCIMOM update](#)

- tcl
- tk
- expect
- conserver-8.1

[Back to top](#)

AIX Licensed Products (LP)

The following tables address Licensed Product (LP) components of AIX relevant to High Performance Computing (HPC) and Clusters software in this service pack. You need to know which releases of the LP components are compatible with your AIX release, as well as the latest fileset fix levels available at the time this HPS Service Pack was released.

AIX 5.3 and AIX 5.2 Licensed Products (LPs)

The following table lists LPs for AIX 5.3 and 5.2 and indicates which LP releases are compatible with which AIX version.

LP	AIX 5.3	AIX 5.2
VSD	4.1.0	4.1.0
LAPI	2.4.4	2.3.3
HPS/SNI (Devices)	1.2.0	1.1.3
PPE	4.3.1	4.2.2
	4.2.2	4.1.1 (*EOS 4/30/07)
	4.1.1 (*EOS 4/30/07)	
LoadLeveler	3.4.2	3.3.2
	3.3.2	3.2.0 (*EOS 4/30/07)
GPFS	3.2.0	3.1.0
	3.1.0	2.3.0
	2.3.0	2.2.1 (*EOS 4/30/2007)
		2.1.0 (*EOS 4/1/2006)
CSM	1.7.0	1.5.1
	1.5.1	1.4.1 (*EOS 4/30/07)
	1.4.1 (*EOS 4/30/07)	
RSCT	2.4.8	2.3.12
ESSL	4.3.0	4.2.0
Parallel ESSL	3.3.0	3.2.0
	3.2.0	

*EOS - Indicates that this LP release has reached its End of Service date and is no longer supported by IBM.

Detailed LP level check

The LP fileset fix levels listed in the following table indicate the latest available levels at the time this HPS Service Pack was released. They are available from the following website:

[Fix Central: System p support fixes](#)

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

AIX 5.3: [APAR IZ10243](#)

AIX 5.2: [APAR IZ10243](#)

To order the latest APAR for an individual LP component, find the LP name in the table, and then click the APAR link that corresponds to your Release of AIX. The link opens the Packaging Options page in Fix Central, where you select your OS level, and then Continue to the download page.

Alternatively, clicking on the aforementioned Order ALL Fixes 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.

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

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

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

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

LP	Release	Component ID	APAR	Level check	Info
VSD	410	5765G2602	AIX 5.3 IZ10232	rsct.vsd.cmds 4.1.0.18	View
			AIX 5.2 IZ00329	rsct.vsd.rvsd 4.1.0.17	
				rsct.vsd.vsdd 4.1.0.22	
				rsct.vsd.vsdrm 4.1.0.7	
LAPI	244	5765G2601	AIX 5.3 IZ10233	rsct.lapi.nam 2.4.4.0	View
				rsct.lapi.rte 2.4.4.4	
				rsct.lapi.samp 2.4.4.0	

	233	5765G2601	AIX 5.2 IZ08271	rsct.lapi.nam 2.3.3.3 rsct.lapi.rte 3.3.3.9 rsct.lapi.samp 2.3.3.0	View
HPS/SNI (Devices)	120	5765G2400	AIX 5.3 IZ10234	devices.chrp.IBM.HPS.rte 1.2.0.11 devices.common.IBM.sni.ml 1.2.0.2 devices.common.IBM.sni.ntbl 1.2.0.1 devices.common.IBM.sni.rte 1.2.0.10 devices.msg.en_US.chrp.IBM. HPS.rte 1.2.0.2 devices.msg.en_US.common. IBM. sni.ml 1.2.0.0 devices.msg.en_US.common. IBM. sni.ntbl 1.2.0.1	View
	113	5765G2400	AIX 5.2 IZ05133	devices.chrp.IBM.HPS.rte 1.1.3.13 devices.common.IBM.sni.ml 1.1.3.2 devices.common.IBM.sni.ntbl 1.1.3.1 devices.common.IBM.sni.rte 1.1.3.11 devices.msg.en_US.chrp.IBM. HPS.rte 1.1.3.1 devices.msg.en_US.common. IBM. sni.ml 1.1.3.1 devices.msg.en_US.common. IBM. sni.ntbl 1.1.3.1	View

431	5765F8300	AIX 5.3 IZ10236	ppe.man 4.3.1.2 ppe.perf 4.3.1.2 ppe.poe 4.3.1.5 ppe.pvt 4.3.1.2	View
-----	-----------	---------------------------------	---	----------------------

PPE

422	5765F8300	AIX 5.3 IZ10235 AIX 5.2 IZ10235	ppe.man 4.2.2.2 ppe.perf 4.2.2.3 ppe.poe 4.2.2.14 ppe.pvt 4.2.2.0	View
-----	-----------	--	--	----------------------

342	5765E6900	AIX 5.3 IZ10238	LoadL.full 3.4.2.1 LoadL.msg.en_US 3.4.2.1 LoadL.so 3.4.2.1 LoadL.webui 3.4.2.1	View
-----	-----------	---------------------------------	--	----------------------

LoadLeveler

332	5765E6900	AIX 5.3 IZ10237 AIX 5.2 IZ00336	LoadL.full 3.3.2.14 LoadL.msg.En_US 3.3.2.5 LoadL.msg.en_US 3.3.2.5 LoadL.so 3.3.2.14	View
-----	-----------	--	--	----------------------

320	5765G66AP	AIX 5.3 IZ10241	gpfs.base 3.2.0.2 gpfs.docs.data 3.2.0.2 gpfs.msg.en_US 3.2.0.2	View
-----	-----------	---------------------------------	---	----------------------

GPFS

310	5765G66AP	AIX 5.3 IZ10240 AIX 5.2 IZ10240	gpfs.base 3.1.0.16 gpfs.docs.data 3.1.0.5 gpfs.msg.en_US 3.1.0.12	View
-----	-----------	--	---	----------------------

230	5765F64AP	AIX 5.3 IZ10239 AIX 5.2 IZ00338	gpfs.base 2.3.0.26 gpfs.docs.data 2.3.0.8 gpfs.msg.en_US 2.3.0.14	View
-----	-----------	--	---	----------------------

CSM	1.7.0	5765F6700	AIX 5.3 IZ08269	csm.bluegene 1.7.0.0 csm.client 1.7.0.1 csm.core 1.7.0.1 csm.deploy 1.7.0.1 csm.diagnostics 1.7.0.1 csm.dsh 1.7.0.1 csm.essl 1.7.0.1 csm.gpfs 1.7.0.1 csm.gui.dcem 1.7.0.0 csm.gui.websm 1.7.0.0 csm.hams 1.7.0.0 csm.hc_utils 1.7.0.1 csm.hpsnm 1.4.1.18 csm.ll 1.7.0.1 csm.pe 1.7.0.1 csm.pessl 1.7.0.1 csm.server 1.7.0.1 (install only on CSM Management Servers)	View
				RPMS:	
				expect 5.42.1-3 openCIMOM 0.8-1 openssl 0.9.7d-2 AIX-rpm 5.3.7.0-7 tcl 8.4.7-3 tk 8.4.7-3 conserver 8.1.7-2	
				csm.server 1.6.0.13 csm.pessl 1.6.0.11 csm.pe 1.6.0.10 csm.ll 1.6.0.10 csm.hc_utils 1.6.0.11 csm.hams 1.6.0.10 csm.gui.websm 1.6.0.10 csm.gui.dcem 1.6.0.10 csm.gpfs 1.6.0.10 csm.essl 1.6.0.11 csm.dsh 1.6.0.11 csm.diagnostics 1.6.0.10	
	160	5765F6700	AIX 5.2 IZ10956	csm.deploy 1.6.0.11 csm.core 1.6.0.12	View

			csm.client 1.6.0.11 csm.bluegene 1.6.0.10 RPMS: expect 5.32-1 openCIMOM 0.8-1 openssl 0.9.7d-2 AIX-rpm 5.2.0.95-1 tcl 8.3.3-1 tk 8.3.3-1 conserver 8.1.7-2	
			csm.bluegene 1.5.1.1 csm.client 1.5.1.3 csm.core 1.5.1.4 csm.deploy 1.5.1.4 csm.diagnostics 1.5.1.0 csm.dsh 1.5.1.6 csm.essl 1.5.1.0 csm.gpfs 1.5.1.0 csm.gui.dcem 1.5.1.0 csm.gui.websm 1.5.1.0 csm.hams 1.5.1.1 csm.hpsnm 1.4.1.17 csm.ll 1.5.1.0 csm.pe 1.5.1.0 csm.pessl 1.5.1.0 csm.server 1.5.1.6 RPMS: expect 5.32-1 openCIMOM 0.8-1 openssl 0.9.7d-2 AIX-rpm 5.2.0.95-1 tcl 8.3.3-1 tk 8.3.3-1 conserver 8.1.7-2	
151	5765F6700	AIX 5.3 IZ05141 AIX 5.2 IZ05141		View

			devices.chrp.base.ServiceRM 1.3.0.70 rsct.basic.hacmp 2.4.8.0 rsct.basic.rte 2.4.8.1 rsct.basic.sp 2.4.8.0 rsct.compat.basic.hacmp 2.4.8.0 rsct.compat.basic.rte 2.4.8.0 rsct.compat.basic.sp 2.4.8.0 rsct.compat.clients.hacmp 2.4.8.0 rsct.compat.clients.rte 2.4.8.0 rsct.compat.clients.sp 2.4.8.0 rsct.core.auditrm 2.4.8.1 rsct.core.errm 2.4.8.1 rsct.core.fsrn 2.4.8.0 rsct.core.gui 2.4.8.0 rsct.core.hostrn 2.4.8.0 rsct.core.lprn 2.4.8.0 rsct.core.rmc 2.4.8.1 rsct.core.sec 2.4.8.1 rsct.core.sensorrm 2.4.8.1 rsct.core.sr 2.4.8.0 rsct.core.utils 2.4.8.1 rsct.opt.saf.amf 2.4.8.0 rsct.opt.storagerm 2.4.8.1	
248	5765F07AP	AIX 5.3 IZ08274		View

RSCT

			devices.chrp.base.ServiceRM 1.2.0.116 rsct.basic.hacmp 2.3.12.0 rsct.basic.rte 2.3.12.1 rsct.basic.sp 2.3.12.0 rsct.compat.basic.hacmp 2.3.12.0 rsct.compat.basic.rte 2.3.12.0 rsct.compat.basic.sp 2.3.12.0 rsct.compat.clients.hacmp 2.3.12.0 rsct.compat.clients.rte 2.3.12.0 rsct.compat.clients.sp 2.3.12.0 rsct.core.auditrm 2.3.12.1 rsct.core.errm 2.3.12.1 rsct.core.fsrn 2.3.12.0 rsct.core.gui 2.3.12.0	
2.3.12	5765F07AP	AIX 5.2 IZ08272		View

rsct.core.hostrm 2.3.12.0
 rsct.core.lprm 2.3.12.0
 rsct.core.rmc 2.3.12.1
 rsct.core.sec 2.3.12.1
 rsct.core.sensorrm 2.3.12.1
 rsct.core.sr 2.3.12.0
 rsct.core.utils 2.3.12.1
 rsct.opt.saf.amf 2.3.12.0
 rsct.opt.storagerm 2.3.12.1

330	5765F8400	AIX 5.3	<p>pessl.rte.common 3.3.0.0 pessl.rte.hv 3.3.0.0 pessl.rte.rs1 3.3.0.0 pessl.rte.smp 3.3.0.0 pessl.man.en_US 3.3.0.0 pessl.rte.mp 3.3.0.0 pessl.rte.rs2 3.3.0.0 pessl.rte.up 3.3.0.0 pessl.loc.license 3.3.0.0</p>	View
-----	-----------	---------	--	----------------------

Parallel ESSL

320	5765F8400	AIX 5.3 PK21664 AIX 5.2 PK21664	<p>pessl.rte.common 3.2.0.1 pessl.rte.hv 3.2.0.1 pessl.rte.rs1 3.2.0.1 pessl.rte.smp 3.2.0.1 pessl.man.en_US 3.2.0.1 pessl.rte.mp 3.2.0.0 pessl.rte.rs2 3.2.0.0 pessl.rte.up 3.2.0.0 pessl.loc.license 3.2.0.0</p>	View
-----	-----------	---	--	----------------------

430	5765F8200	AIX 5.3 PK52484 AIX 5.2 PK52484	<p>essl.msg.en_US.4.3.0.0 essl.msg.En_US.4.3.0.0 essl.rte.common.4.3.0.2 essl.rte.rs1.4.3.0.2 essl.rte.smp.4.3.0.2</p>	View
-----	-----------	--	--	----------------------

ESSL

420

5765F8200

AIX 5.3

[PK20638](#)

AIX 5.2

[PK20638](#)

essl.rte.common 4.2.0.5

essl.rte.rs1 4.2.0.5

essl.rte.rs2 4.2.0.0

essl.rte.smp 4.2.0.5

essl.rte.mp 4.2.0.0

essl.rte.up 4.2.0.0

essl.man.en_US 4.2.0.0

essl.loc.license 4.2.0.0

essl.msg.en_US.4.2.0.1

essl.msg.En_US.4.2.0.1

[View](#)

[Back to top](#)

Health check/planning help

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

- [Cluster support](#)
- [HPS](#)

IBM HPS POWER 4 Service Pack 22 Addendum 1 FLASH/Readme

Installation Sequence

Anchor link navigation

- [Introduction](#)
- [Supporting documentation](#)
- [Software updates](#)
- [Component updates/download information](#)

- **[Installation sequence](#)**
- [Installation guidelines](#)
- [Known problems and workarounds](#)
- [Fix list and readme archive](#)

Recommended installation sequence (overview)

The following overview lists the recommended installation sequence for installing software, firmware and operating system updates.

Anchor link navigation

- [HMC](#)
- [SNM/FNM](#)
- [HPS/SNI](#)
- [GFW](#)

- [Power Subsystem Microcode](#)
- [AIX](#)
- [AIX LPPs](#)

1. HMC software

Important Preliminary steps

- Verify that the HMC is at HMC V3R3.0 or higher before you install this update (required).
- Check if a BIOS update is required on the HMC.
- Check if BIOS hyperthreading is to be disabled on the HMC.
- Disable the SNM/FNM software from the Switch Network Management Panel.

Perform one of the following installation tasks:

- New Install HMC from Recovery CD *or*
-
- Install Upgrade from Recovery CD *or*
-
- Update from .zip file (web)
-

Level Check: Verify that the HMC Code Level is the [current level](#) shown in [Component update/download information](#).

2. SNM software

For detailed information, link to the specified SNM ReadMe from the HMC page under the SNM tab.

For a thorough Discussion of the SNM GUI Control panel , please see: [Appendix E. Switch Network Manager \(SNM\) in pSeries High Performance Switch Planning, Installation, and Service](#).

Level Check: Verify that the SNM Code Level is the **current level** shown in [Installation guidelines, SNM Component update/download information](#).

3. HPS/SNI LP software

Important:

If you are upgrading from Service Pack 6 or below - do not reboot logical partitions (LPARs) until after Step 4 Install GFW is complete!

See Problem #1 in the "Known Problems" section of this document for more information.

Level Check: Verify that the HPS Code Level is the current level as shown in the [Component update/download information](#).

4. GFW software

Using the recommended AIX command line (update_flash) method with a locally available GFW img file.

Level Check: Verify that the GFW Code Level is the [Component update/download information](#).

For detailed download and unpacking procedures, see the [FAQ for Microcode downloads](#).

5. Install Power Subsystem Microcode on each frame

Level Check: Verify that the Power Code Level is the current level shown in [Component update/download information](#).

6. Install AIX base updates and any PTFs on each node

Level Check: See [AIX download info](#) for the recommended AIX service level for this Service Pack.

7. Install AIX LP updates on each node

Note: CSM LP's need to download and install the openCIMOM-0.8-1 RPM 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

Level Check: see "[Detailed LP Level Check](#)"

[Back to top](#)

Health check/planning help

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

- [Cluster support](#)
- [HPS](#)

IBM HPS POWER 4 Service Pack 22 Addendum 1 FLASH/Readme

Installation Guidelines

Anchor link navigation

- [Introduction](#)
- [Supporting documentation](#)
- [Software updates](#)
- [Component updates/download information](#)

- [Installation sequence](#)
- **[Installation guidelines](#)**
- [Known problems and workarounds](#)
- [Fix list and readme archive](#)

Overview

Install the software in the order listed in this section.

Anchor link navigation

- [HMC](#)
- [SNM/FNM](#)
- [HPS/SNI](#)
- [GFW](#)

- [Power Subsystem Microcode](#)
- [AIX](#)
- [AIX LPPs](#)

HMC

Installing HMC software consists of performing some important preliminary steps, and then choosing one of the three available methods for installing the HMC software itself.

[Important preliminary steps](#)

[HMC installation procedures](#)

Perform one of the following installation tasks:

- [New Installation of the HMC from the Recovery CD](#) OR
- [Installation Upgrade from the Recovery CD](#) OR
- [Update from .zip file](#) (web)

The HMC installed with this Service Pack has an updated HMC WebSM PC CLIENT. To reinstall this version:

Uninstall - reinstall HMC WebSM PC CLIENT (Install Shield version)

Level Check: Verify that the HMC Code Level is the current level shown in [Component updates/download information](#) section.

Important preliminary steps

1. Verify HMC Code level.

Installation of HMC Recovery CD requires an upgrade install if you are installing an existing HMC whose version is R3 V2.6 or less.

Please read [HMC corrective service](#) for important information regarding HMC Recovery CDs. Contact your IBM Sales Representative or Business Partner to order HMC Machine Specific CDs.

2. Check if BIOS update is required on HMC. There is a mandatory BIOS upgrade required for these HMC PC's:

7315-C03, 7310-C03, 7315-CR2, 7310-CR2

If your HMC model is not listed, skip this step. If you are updating the HMC on a listed model, then you must first update the BIOS of that HMC model. The BIOS and install instructions can be obtained by linking to the "BIOS Updates" for the referenced machine from the [HMC POWER4 servers page](#). This BIOS will also ship as part of Feature Code 0960.

3. Check if BIOS hyperthreading is to be disabled on HMC.

Many of the rack mounted HMC's (8187-KUH, 7315-C03) have a BIOS option to enable hyperthreads. The imbedded kernel will not run well when this option is enabled. You must disable this setting before upgrading to HMC3.3.5.

4. Decide whether to install, upgrade or update.

Beginning with HMC Version 3 Release 3.2, (Service Pack 10) the Ext3 (JFS) filesystem will be enabled if customers [perform a New Install/Upgrade](#) to this new level of code by using the HMC Recovery CDs. The Ext3 filesystem is a journaled filesystem and is more reliable and less prone to corruption in case of unexpected loss of power on the HMC.

Please note the difference between Upgrade and Update:

- **Upgrade** is done via cd load as described in [Perform a New Install/Upgrade](#).
- **Update** is done via downloading a zip file as described in [Perform an Update](#).

Notes: Updating to HMC Version 3 Release 3.3 or later using the Install Corrective Service will not enable this feature.

For an Upgrade installation, the following steps must be taken before rebooting the HMC for the Upgrade process:

1. Ensure that the user's home directories are not filled up with debug data. The Upgrade partition only has 2GB in free space to preserve the upgrade data.
2. Ensure that HMC debug is turned off by running the **pedbg -d off** command.

With debug enabled, certain log files are locked for writing and prevent the Save Upgrade Data task from completing.

3. Perform the Save Upgrade Data task from the Software Maintenance Panel on the HMC console. This task should be run immediately before rebooting the HMC with volume 1 of the recovery CD. If the HMC reboot does not go to the install menu of volume 1 of the recovery CD, you should repeat the save upgrade task.

The procedure for both Installation and Upgrade is identical except:

- For New Installation: When asked to perform an Install/Recovery or Upgrade, select Install/Recovery F8.
- For Upgrade Installation: When asked to perform an Install/Recovery or Upgrade, select Upgrade F1.

Installation steps

1. Disable the SNM/FNM software from the Switch Network Management Panel.

Disable SNM Software on **ALL** HMC's attached to the cluster. For each HMC attached to the

cluster, select the "Disable SNM Software" option to open a task dialog box that stops the SNM daemon for both verification mode and normal mode.

Refer to [Appendix E. Switch Network Manager \(SNM\)](#) in pSeries High Performance Switch Planning, Installation, and Service.

2. Perform a New Install/Upgrade:

The HMC Recovery CD package is now a set of two CDs. Complete the following steps to perform the upgrade:

[HMC Recover CD Version 3 Release 3.5](#)

- a. Reboot the HMC with volume 1 of the recovery CD inserted in the DVD Ram drive.

Note: If the HMC fails to boot volume 1 of the recovery CD, the boot sequence in the HMC BIOS may need to be changed so that the DVD/CDROM is before the hard disk in the startup sequence. If you have run the save upgrade data task before the startup sequence was set correctly, then you should rerun the save upgrade data task before installing the HMC with volume 1 of the recovery CD.

- b. Select F8 for New Installation

OR

Select F1 for an Upgrade installation. (**Note:** This is *not* the same as an UPDATE)

- c. On the next screen, select F1 to confirm your selection.

The Install/Upgrade process proceeds until you are prompted to insert the second CD.

- d. Remove the CD from the DVD Ram drive and hit enter when the install is completed.

3. Perform an Update.

If you are *updating* from Service Pack 7 (HMC V3 R3.0) or higher and choose to do the update, then install the following HMC PTF:

HMC_Update_V3R3.7.zip

You can download this update from this location: [HMC Version 3.3.X corrective service](#).

Install from the HMC support link ONLY if the HMC is at a Release 3 Version 3.X level (3.3.X):

1. Select the **HMC_Update_V3R3.7.zip** link to download the update package.
2. Use the **HMC --> Install Corrective Service** option to install.
3. Reboot HMC after a successful installation.

You may install this update directly from the web via "HMC Install Corrective Service". See Step 3, Install Update, in the [HMC update for V3R3.7 Readme](#).

At the HMC interface, follow these steps to install the update:

1. Select **Software Maintenance**.
2. Select **Install Corrective Service**.
3. Select **HMC**.
4. Select **Download corrective service from remote system**.
5. Enter the specified information in the following fields:

Remote Site: ftp.software.ibm.com

Patch File: /software/server/hmc/updates/HMC_Update_V3R3.7.zip

User ID: anonymous

Password: <your email address>

The HMC interface retrieves the update package from the remote FTP server and begins the install process.

Reboot the HMC after the installation of the update has completed. Rebooting ensures that all changes are available immediately.

After the HMC is rebooted, to Verify a successful update:

1. Select Help in the top menu bar.
2. Select About the Hardware Maintenance Console for pSeries.
3. On the "About" splash panel, check for the following information:
The Version is 3
The Release is 3.7

[Back to top](#)

Web-based System Manager Remote Client (WebSM PC Client)

After the HMC Install New/Upgrade or Update is complete, the Install Shield version of the Web-based System Manager Remote Client (WebSM PC CLIENT) may need to be reinstalled on your remote server or PC. Follow these instructions:

1. Uninstall an existing WebSM client:

Use the "Add/Remove Programs" option on the Windows Control Panel. Select "Web-Based System Manager Remote Client, and then click the Change/Remove button.

2. Install a new WebSM client:

Use your preferred web browser to visit the recently installed HMC at **http://<hmc-hostname>/remote_client.html**. This download and the installation may take ten to twenty minutes.

For pertinent information see the [WebSM discussion](#) in the Software updates section of this Readme.

For complete details refer to Chapter 9, Installing and Using the Remote Client, in the [Hardware Management Console for pSeries Installation and Operations Guide](#). (SA38-0590-07)

HMC level check

Check the HMC level as follows:

On the command line:

The output of the **lshmc -V** shows:

Version: 3

Release: 3.7

HMC Build level 20060207.1

On the GUI:

Displaying the splash panel by selecting **Help > About Hardware Management Console** from the menu shows:

Version: 3

Release: 3.7

HMC Build level 20060207.1

HMC Important Notes

Install the HMC code by following the instructions in the HPS Guide. Have your IBM CE download the most recent copy of the HPS guide from IBM CORE to get updated HPS install information. You should also review the HMC information on the web page where you downloaded the images.

This Service Pack **REQUIRES**:

HMC V3.3.0 (required since Service Pack 6 release)

- HMC V3.3.0 is a NEW BASE release of the HMC introduced in Service Pack 6 that uses a new imbedded kernel.

- This version MAY OVERWRITE root directories (for /, /home/root and /home/hscroot) deleting any scripts that may be there.
- This version MAY delete the Power Subsystem Microcode RPM images on your HMC and you will have to reacquire it for future installs.

The HMC is now installed using two CDs.

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

New Installation and PTF update installation are supported for this release.

Upgrade installation is only supported when upgrading from HMC 3.2.X or greater.

As part of any system change, it is recommended that you maintain a hard copy of network connections, 8 port/ran box configurations and Switch Group IP's.

- Network connections are on the GUI:
HMC Maintenance Panel => System Configuration => Customize Network Settings:
IP Address and Netmask for Ethernet0 and Ethernet1, Default Gateway, Nameserver, Domain
- 8 port RAN box configurations are on the GUI:
HMC Maintenance Panel => System Configuration => Configure Serial Adapter:
Option 2 shows the current configuration.
- Switch Group IP's are on the GUI: Switch Management => Switch Utilities => Switch Group Configuration

Known problems and issues with the HMC V3 R3.0 Environment

- wu-ftp will be removed from the HMC distribution. One will be able to ftp out of the HMC but not into the HMC. The 'scp' command is available if you enable secure shell (ssh).
- The websm PC client has a performance decline when downloading the plugin classes from the server. The first time an operation is performed using the client, the task may be slow to launch. Subsequent use of the task, will respond as normal.

Retain Tip on how to use pesh

To give IBM support personnel the ability to retrieve certain trace/debug information on the HMC, the customer should create a user "hscpe" and assign a password. IBM support can contact the customer to get the password, and then remotely connect to the HMC (with customer consent).

This allows IBM support to perform additional functions, such as viewing logs or starting trace to

diagnose problems on the HMC. This user has access similar to the "hscroot" user on HMC. When accessing the HMC remotely via ssh, the "hscpe" user is put into the restricted shell environment. To bypass the restricted shell, the pesh command is provided. The pesh command can only be run by the "hscpe" user, allowing this user to pass in the serial number of the HMC. If the serial number is correct, the user is required to enter a password obtained from IBM Support. If the password is correct, then the user is then put into the un-restricted shell as user "hscpe".

This allows IBM support to perform additional functions, such as viewing logs or starting trace to diagnose problems on the HMC. This user has access similar to the "hscroot" user on HMC. When accessing the HMC remotely via ssh, the "hscpe" user is put into the restricted shell environment. To bypass the restricted shell, the pesh command is provided. The pesh command can only be run by the "hscpe" user, allowing this user to pass in the serial number of the HMC. If the serial number is correct, the user is required to enter a password obtained from IBM Support. If the password is correct, then the user is then put into the un-restricted shell as user "hscpe".

Example:

```
pesh 23A345K
```

Enter the serial number in upper case letters.

You will be prompted for a password. Enter password that was provided by IBM Support in lower case letters.

The HMC serial number can be queried using the command, "lshmc -v | grep SE" or read from the label that is on the front of the HMC. Use the command "date" to verify that the date of the HMC is for the day you intend to use the pesh command.

Starting with HMC Version 3 Release 3.0 and Version 4 Release 1.0, user can also access the restricted shell terminal on the local HMC, by right mouse click on the desktop and selecting the Terminal--rshterm task. If one login at the HMC as user hscpe, the pesh command can also be run from the restricted shell terminal.

For HMC Version 3 Release 3.0 and below, the "hscpe" user id can be created with any role, however, in order to use some of the High Performance Switch (HPS) debug commands, the Service Rep role needs to be selected.

For new HMC installation(s) follow the instructions as described in IBM Hardware Management Console for pSeries Installation and Operations Guide.

To understand how to connect the rs422/rs232 cables see the HPS Guide: Chapter 6; Step 6, "Install the Hardware Management Console (HMC)" thru Step 16. "Verify Installation is Complete"

For "Code load requirements for existing server frames" see Chapter 6

For p655 "Code load requirements for existing p655 server frames"

For p690 "Code load requirements for existing p690 server frames"

[Back to top](#)

SNM/FNM Installation Guidelines

Install SNM Software

Follow these steps to install SNM software:

1. Disable SNM Software on ALL HMCs attached to the cluster. The SNM/FNM software should already be disabled as part of the installation procedure for HMC software. If the SNM software is not disabled, disable it now.
2. Install SNM update, required on all of the HMCs in the cluster. (The version of SNM distributed with the HMC code is not the most current version.)

To install the corrective service directly from the Internet, use the HMC Corrective Service GUI as follows:

Remote Site: **ftp.software.ibm.com**

Patch File: **/software/server/hmc/fixes/<SNM Update zip file>**

User ID: **anonymous**

Password: **<your email address>**

The HMC GUI retrieves the update package and begins the install process.

3. Reboot all HMCs to complete SNM software update.

Note: Do not reboot HMCs until the corrective service has been successfully installed on all HMCs.

After the HMC is rebooted verify that the corrective service update was successful.

DO NOT ENABLE SNM SOFTWARE AT THIS TIME

Level Check:

Select Switch Network Management from the Switch Management folder in the Navigation area. The Status line, the last line on the Switch Network Management panel, should show the [Current STATUS RPM Version](#).

Notes:

Refer to the [pSeries High Performance Switch \(HPS\) Planning, Installation and Service Guide \(GA22-7951-01\) \(HPS Guide\)](#) for more details on the Switch Network Manager. Review the following chapters:

- o Chapter 4. System management components
- o Chapter 6. Installation: Bringing the network online: Step 1: Enable SNM
- o Appendix E. Switch Network Manager (SNM) - The SNM Graphical User Interface

The SNM GUI does NOT update its view automatically. You MUST refresh the display via one of the following:

- o The GUI menu "Reload" button
- o The "Menu">"View">"Reload" function
- o The 'F5' key

HPS/SNI LP Installation Guidelines

Install HPS/SNI LP Software.

1. Application of AIX provides base level HPS. Apply HPS/SNI LP Base fileset images to LPARs. (See [AIX Licensed Products](#) in the **Component updates** section.

Leave SNM Software disabled. To verify the service levels for HPS/SNI LPs, on each LPAR issue:

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

or

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

Refer to Section 3 - Detailed LP Level Check for correct levels.

If you are upgrading from Service Pack 6 or lower **DO NOT** reboot logical partitions (LPARs). See Note 1. Go to Step 2.

If you are upgrading from Service Pack 7 or higher, verify that the 64-bit kernel is currently in use. Verify the TLP settings. See Note 2.

Reboot LPARs. You should reboot the LPARs as soon as possible to properly integrate the changes and to avoid disruption of current functionality. Go to the [Install GFW](#) procedure.

Notes:

1. If you are upgrading from Service Pack 6 or below, do not reboot logical partitions (LPARs) until [GFW installation](#) is complete. Rebooting prematurely will generate "phantom" SNI devices. See

the **Known problems and workarounds** section for more information.

2. If you are upgrading from Service Pack 7 or higher you should already have set up LPARs to boot the 64 bit kernel and enabled the Technical Large Page (TLP) option, as described in the next step.

Verify 64-bit kernel is currently in use (on an LPAR): `bootinfo -K 64`

Verify the TLP settings:

```
vmo -a | grep lg
```

```
lgpg_size = 16777216
```

```
lgpg_regions=YYY where YYY is the number of technical large pages to export
```

```
(Ex: lgpg_regions = 256)
```

```
soft_min_lgpgs_vmpool = 0
```

3. It is strongly recommended that users read and become familiar with the items covered in [HPS/SNI Notes](#).
4. IBM pSeries HPS now requires that you set up LPARs with 64-bit kernel and enable Technical Large Page (TLP) option.

After successful installation of HPS Filesets from levels at Service Pack 6 or lower, 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
2. `ln -fs /usr/lib/boot/unix_64 /unix`
3. `ln -fs /usr/lib/boot/unix_64 /usr/lib/boot/unix`
4. Determine which rootvg hdisk contains the boot logical volume (usually hd5). This hdisk will be your "ipldevice".
 - 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..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`

Regarding the use of memory by TLP

Also refer to [Regarding the use of memory by RDMA and TLP](#) in the **Software updates** section.

To set up Large Page Option:

For configuration details, see Large Page Support in the following publication:

[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.

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

```
vmo -r -o v_pinshm=1 -o lpgg_size=16777216 -o lpgg_regions=YYY
```

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

For Example:

To set up a node with 8 sni adapters:

16MB Large Page: lpgg_size = 16777216

256 Large Pages: lpgg_regions = 256

```
dsh <nodelist> "echo y|vmo -r -o v_pinshm=1 -o lpgg_size=16777216 -o lpgg_regions=256"
```

Use "echo y|vmo", otherwise vmo will prompt for verification to run bosboot)

Would generate this response:

```
Setting v_pinshm to 1 in nextboot file
```

```
Setting lpgg_size to 16777216 in nextboot file
```

```
Setting lpgg_regions to 256 in nextboot file
```

```
Warning: some changes will take effect only after a bosboot and a
reboot
```

```
Run bosboot now?
```

```
bosboot: Boot image is 19624 512 byte blocks.
```

```
Warning: changes will take effect only at next reboot
```

Note: The vmtune sample program is being phased out and is not supported in future releases. It is replaced with the vmo command (for all the pure VMM parameters) and the ioo command (for all the I/O related parameters) which can be used to set most of the parameters that were previously set by vmtune. The -v flag has been added to vmstat to

replace the -A flag which display counter values instead of tuning parameters. For AIX 5.2, a compatibility script calling vmo and ioo is provided to help the transition.

To Check that Large Page Option is set:

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

HPS/SNI Notes

Notes on using TLP (Large Page) Settings in an HPC environment

It is strongly recommended that users familiarize themselves with TLP basics and configuration options available to them, at this location. Federation switch adapter requires TLP usage and these TLP requirements are documented (see "Here is a formula to calculate the required TLP" in the HPS/SNI LPs section below) in a latter section of this document.

You should also consult the section on Large page feature on AIX in the following publication:

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

Note: Users need to be aware of the usage of the LoadLeveler pre-emption features with TLP (Large Pages), specifically 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 you want LoadLeveler to schedule jobs based on the availability of large page, (especially if the job is going to run in mandatory Large Page mode) you may consider using the LoadLeveler consumable resource feature. The function has been available for several years and is documented in the LoadLeveler manual.

Notes on tuning Virtual Memory 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 cluster size, the amount of system memory and the number of processors per CEC. When tuned properly, these settings will ensure enough memory remains available for core cluster infrastructure applications (RSCT, GPFS, LL). The recommended initial value for these tunables are:

```
minfree = 10000
```

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

Note: Tuning these settings helps you 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.

[Back to top](#)

GFW Installation Guidelines

The recommended installation is the AIX command line method - `update_flash` - using locally available GFW img file.

Important Preliminary Notes

Notes on updating GFW code (system firmware) from the AIX command line

Document Reference: pSeries High Performance Switch Planning, Installation, and Service.

For "Code load requirements for existing p690 and p655 server frames" please see the section in Chapter 6 titled "Step 3: p690 GFW code load" or "Step 3: p655 GFW code load" respectively in the HPS guide.

To understand how to connect the rs422/rs232 cables see the HPS Guide:

Chapter 6; Step 6, "Install the Hardware Management Console (HMC)" thru Step 16. "Verify Installation is Complete"

For "Code load requirements for existing server frames" see Chapter 6

For p655 "Code load requirements for existing p655 server frames"

For p690 "Code load requirements for existing p690 server frames"

For each CEC on which you want to install the GFW code, one partition running AIX must have

"Service Authority" set. Linux does not support microcode download at this time.

The "Service Authority" is set on one LPAR per CEC in the LPAR's profile "other" tab. This designates the LPAR as authorized to provide update images to the CSP. All partitions except the one with "Service Authority" must be shut down.

The partition with "Service Authority" must own the device from which the microcode update image will be read. It is also recommended that the partition with "Service Authority" have a hard disk.

If the required devices are not in the partition with "Service Authority", the customer or system administrator must reassign the appropriate resources to it. This requires rebooting the partition with "Service Authority".

If the firmware on a full system partition is being updated, no special steps are required to perform the firmware update using the service aid.

Ensure the GFW image file is not corrupted/truncated before you begin the update_flash process.

Check that /var and /tmp directories are not above 50% full on the partition with the service authority.

The update process can range from 20 minutes to 2 hours, depending on system configuration. The system reboots itself during the update process. Since SNM is disabled during this process, the SNI adapter interfaces will NOT be configured, or will be incorrectly configured when the LPAR(s) reactivate.

It is recommended that you use the ``update_flash -f`` command as opposed to the shutdown -Fu method.

AIX52 APAR IY49146 is required for update_flash to work correctly. Level Check by running this command on the partitions:

```
instfix -ik IY49146
```

The ``update_flash`` command will reboot the CEC(s) and will activate the LPAR(s).

You may find some more detailed instructions provided on the website with the latest image: [pSeries 690 and 670 Version 3 Firmware Update](#)

To install GFW update using Diskette method:

For p690 systems follow the instructions in the HPS guide on "Step 3: p690 GFW (system firmware code load)" in chapter 6.

To install GFW update using NIM method:

For a p655 CEC via NIM, follow the HPS guide Chapter 6, "Code load requirements for existing p655 server frames", Step 3. GFW (system firmware) code load.

Installation steps

This procedure employs the recommended AIX command line (`update_flash`) method)

1. SNM Software should be still disabled. Verify from the SNM GUI Panel.
2. For each CEC on which you want to install GFW code, shut down all partitions except the one with service authority.
3. Install the appropriate GFW driver on each CEC to be upgraded: On the AIX partition with Service Authority:
 - a. Copy the GFW firmware update code to /tmp
 - b. Enter the following command:

```
/usr/lpp/diagnostics/bin/update_flash -qf /tmp/<gfw img file>
```

The system will apply the new firmware, reboot, and return to the AIX prompt.

If you use dsh to invoke `update_flash` then use the `-q` flag so it does not put out a prompt.

Example: `dsh /usr/lpp/diagnostics/bin/update_flash -qf /tmp/<GFW img file>`

More conveniently, the '?' represents a single character 'wild card' and will select either /tmp/3H050405.img or /tmp/3J050405.img. Do not have both on the same LPAR.

```
dsh [-N <nodegrp>] "echo `"/usr/lpp/diagnostics/bin/update_flash -qf tmp/<GFW img file>`"|at now"
```

```
job root.1102366157.a at Mon Dec 6 15:49:17 2004
job root.1102366120.a at Mon Dec 6 15:48:40 2004
job root.1102366242.a at Mon Dec 6 15:50:42 2004
```

4. After the LPAR(s) are Running, power OFF the CEC(s) from the GUI or by using CSM `rpower -not` from EPO red switch.
5. After the CEC(s) are powered off, follow these steps to enable the SNM/FNM Software from the GUI:
 - a. Select Switch Network Management from the Switch Management folder in the Navigation area.
 - b. Select Logical Topology on a 'per plane' basis.
Select the number of Planes and select the Logical Topology on a per plane basis (Number of Endpoints on a plane).
Example: 2 Frames and 2 Switches with *no* Switch-to-Switch Links is 2 planes; 1NSB_0ISB_16EP (16 Endpoints).
Example: 2 Frames and 2 Switches with *any* Switch-to-Switch Links is 1 plane; 2NSB_0ISB_32EP (32 Endpoints).

- c. Enable SNM Software for Normal Operation.

OR

Enable SNM Software for Switch Network Verification.

- d. After Enable SNM Software task completes, verify current HMC and SNM version as needed via the Management Properties task:

The HMC version is under the "Version" tab. The SNM version is under the Management tab, "SNM Version" column. This should show the current version as given in [Component updates/download information](#) section.

Notes:

The "Enable SNM Software for Switch Network Verification" task is used during new system setup/ installation or after reconfiguration to initialize and check out the system by NOT removing bad links or routes. The main purpose is to discover cable mis-wires and defective hardware. No Routing is modified or loaded. It is NOT intended to be used in a working environment.

Refer to the HPS Guide for more details on defining the Switch Network Topology: Review - Chapter 4. System management components; Step 1: Enable SNM Review - Appendix E. Switch Network Manager (SNM) - The SNM Graphical User Interface

The SNM GUI *does not* update its view automatically. You *must* refresh the display via one of the following:

- The GUI menu "Reload" button.
- The "Menu">"View">"Reload" function.
- The 'F5' key.
- Power up CEC(s) from the HMC GUI and activate logical partition(s).
- Enable technical large page support - Required for levels greater than Service Pack 6.
Note: Refer to "[HPS/SNI](#)" section for details on technical large page setup.
- Determining the level of firmware on the processor subsystem
Firmware level is indicated as: 3xyymmdd.img; where
x = a firmware designation such as J or H - J = p655, H=p690
yy = year, mm = month, and dd = day of the release *similar to 3J060626*

Check the GFW microcode level from either A VTERM to the main SP Menu. This should show the correct level on the top line.

OR

You can also check the GFW level from the AIX command line on the active LPAR(s):


```
lscfg -vp | grep alter | grep "\.3"
```

You should see:

ROM Level (alterable)..... < GFW img level>

OR

ROM Level (alterable)..... < GFW img level> where < GFW img level> is the level given in [GFW download information](#).

- Determining the level of HPS adapter microcode:
 - a. From an AIX52 partition, issue:


```
/usr/sni/aix52/debugtools/sni_get_ucode_version -l sniz
```
 - b. From an AIX53 partition, issue:


```
/usr/sni/aix53/debugtools/sni_get_ucode_version -l sniz
```

 where z = sni interface number on your system anywhere from 0 thru 7 (Eg. sni0) which can be seen in "netstat -in" output. This should show the current version as given in [GFW download information](#).

Important: You will need to reinstall the GFW update if you neglected to disable SNM during the GFW update or if you added/replaced an HPS adapter. Otherwise, the HPS ucode may not have been applied correctly.

[Back to top](#)

Power Subsystem Microcode Installation Guidelines

Install Power Subsystem Microcode on each frame by following these steps:

1. Download the Power Subsystem Microcode from the [Microcode downloads](#) site to an FTP server.
 - Install via the HMC GUI through the Software Maintenance -> Frame panels.
 - Receive Corrective Service.
 - Install Corrective Service.

If you are upgrading from Service Pack 9 or above, the installation of Power Subsystem Microcode rpm is complete. Proceed to Step 4 to level check installed version.
2. If you are upgrading from Service Pack 8 or lower, the new Power Subsystem Microcode requires the switches to be recycled for the changes to become effective.
 - From the HMC GUI select "Switch Network Management > Switch Topology View" for each switch plane.
 - Select "Selected > Power >Off".
 - Refresh GUI to verify power status.

- Select "Selected >Power >On".
- Refresh GUI to verify power status.
- Repeat this procedure on all switch planes.
- Recycle the SNM daemon using the HMC GUI by selecting Switch Network Management > Disable SNM Software refresh display: Menu > File > Refresh >Enable SNM Software.

Recommended procedure for a complete power-cycle of the cluster ("EPOW")

(Cold Boot Procedure)

In the event that power-cycling the HPS switch boards fails, as a last resort the following procedure is suggested:

For a complete discussion of the subject :

See the pSeries High Performance Switch Planning, Installation, and Service, Chapter 9, Service procedures, Managed system power on and power off (LPAR reboot).

The 10 minute wait period is for switch frames to stabilize.

Below is the suggested procedure:

1. Shutdown all the LPARs.
 2. `rpower -a cec_off`.
 3. Stop `fnmd`.
 4. EPOW off all the CECs and the ISB frame.
 5. Start `fnmd`.
 6. EPOW on the ISB frame and the NSB frames.
 7. Wait 10 minutes, EPOW everyone else on.
 8. Check for flashing lights on the (optical) risers in the ISBs.
 9. `rpower -a cec_on`.
 10. Tail `/var/hsc/log/*Init.log` until that stops.
 11. Run `hps_check.pl`.
 12. Activate all the LPARs.
- Level Check After completing the pcode installation, go to the HMC GUI and verify successful installation:
 - Select: Software Maintenance --> Frame --> Install Corrective Service
 - Verify that the ["Installed Version"](#) matches the version you just installed.

Important: This window may not automatically refresh when installation completes. Manually refresh the window as necessary.

- Verify HPS is functioning. At this point the basic HPS installation is complete. You should now be able to ping over the switch.

A general check is a good ping all script. The HPS Documentation covers this topic in the host based verification tool.

Post Installation Task 2. (pSeries High Performance Switch Planning, Installation, and Service;
"Run the

If SNM Software is not running then :

The "Switch Management" | Switch Network Management | views : Switch Topology

View and End Point View will not be populated. There is a dialog message indicating that you will not be able to ping over the switch.

The hps_check.pl file is not available in a closed box without the root password.

It should show the links as Timed and MPA Available :

```
Lpar Name Lpar# Sni# => Adapter# Csp# Cronus# => Frame Cage
Chip Port : Timed? MPA TOD
c661f1rp02 1 0 0 2 5 1 1 5 2 YES YES SLV
c661f1rp02 1 1 1 3 4 4 4 5 2 YES YES BAK
```

[Back to top](#)

AIX Installation Guidelines

Install AIX base updates on each node by following the steps in this section.

See [AIX download information](#) for the recommended AIX service level for this Service Pack.

For complete instructions on the installation procedures see the [AIX 5L Version 5.2 Installation Guide](#) and install the recommended AIX 5L version service level for this Service Pack.

1. Go to [Quick links for AIX fixes](#), and then click the "Specific fixes" link for your Release of AIX.
2. In the "Search By" drop down box Select "APAR or abstract".
3. In the "Search String" text box enter the specific fix as given in [AIX download information](#), and then click "Go". The search results page displays a text box with the specific fix requested.
4. Select that package, and then click "Add to my download list".
5. You can view by clicking "View my download list".
6. Repeat these steps for any additional fixes or PTFs.
7. After all the indicated fixes are in the download list click "Continue".
8. Specify your "Current level" (use oslevel -r command to determine your AIX level).
9. After "Select a download server", click "Continue", and then follow the instructions for downloading the selected file sets.

Please read the "Memo to Users" on the Download fixes page for installation and fix information. Level check AIX by running this command on the logical partition(s):

```
lslpp -ha bos.mp64
```

Verify that filesets are at or above the given in [AIX download information](#) for the recommended AIX service level for this Service Pack.

AIX LP Installation Guidelines

Install AIX LP updates on each node by following these steps.

Level check: ["Detailed LP Level Check"](#)

Notes: CSM LPs need to download and install openCIMOM-0.8-1 RPM update.

Document Reference:

- AIX 5L Version 5.2 Installation Guide and Reference (SC23-4389-03)
- IBM RSCT: Administration Guide

As noted in the Installation Guide, in order to use CSM LPs, you need to download and install openCIMOM-0.8-1 RPM update. You can download openCIMOM from <http://www-1.ibm.com/servers/aix/products/aixos/linux/download.html>

Select "Package" OpenCIMOM "Version" 0.8 (5.2) (For AIX 5.2) For complete details see: Chapter 4. Installing the management server; Step 6. Download Open Source Software of CSM Guide.

Document Reference:

- IBM Cluster Systems Management for AIX 5L Planning and Installation Guide.
- AIX 5L Version 5.2 Installation Guide and Reference.
- IBM Reliable Scalable Cluster Technology Administration Guide.

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>

Document Reference:

[IBM Parallel Environment for AIX 5L Installation Version 4, Release 1.1 Chapter 1. Introducing PE 4.1.1.](#)

1. Download and install the applicable IBM Virtual Shared Disk(VSD) , LAPI, HPS, PPE, LoadLeveler, GPFS, Parallel ESSL, ESSL, CSM and RSCT PTF updates.

Apply the listed APARs which are needed for this Service Pack.

They are available from [Quick links to AIX fixes](#) .

If you are upgrading to this Service Pack from a service pack that is earlier than SP9, then installing CSM1.4 is required.

The Maintenance package contains CSM 1.4 which requires RPM update openCIMOM 0.8(5.2).

2. To verify that the service levels for your LP's are current for this Service Pack, 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:// /'
```

Check the levels against those listed in [Detailed LP Level Check](#)

3. RSCT Migration Issues:

See [Known problems and workarounds](#), "hagsglsm is not reporting the local switch membership group"

For more information, refer to RSCT Administration Guide; Chapter 3. Creating and Administering an RSCT Peer Domain; Migration.

[Back to top](#)

Health check/planning help

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

- [Cluster support](#)
- [HPS](#)

IBM HPS POWER 4 Service Pack 22 Addendum 1 FLASH/Readme

Known Problems, Work Arounds, and Restrictions

Anchor link navigation

- [Introduction](#)
- [Supporting documentation](#)
- [Software updates](#)
- [Component updates/download information](#)

- [Installation sequence](#)
- [Installation guidelines](#)
- **[Known problems and workarounds](#)**
- [Fix list and readme archive](#)

Overview

Anchor link navigation

- [Known problems and work arounds](#)

- [Restrictions](#)

The following known issues and restrictions exist for this Service Pack

Known problems and workarounds

- [The new AIX 5.3 TL06 \(kernel filesets 5.3.0.60\) has a change in the pthread scheduling that breaks GPFS synchronization.](#)

- [Phantom SNI devices may appear after upgrade from pre-Service Pack 7.](#)
- [hagsglsm is not reporting the local switch membership group.](#)

The new AIX 5.3 TL06 (kernel filesets 5.3.0.60) has a change in the pthread scheduling that breaks GPFS synchronization.

Users Affected:

Users running AIX 5.3 TL06 or higher.

Problem Description:

Unpredictable results can occur with the GPFS subsystem due to problems with GPFS synchronization mechanisms.

Recovery Procedure:

Contact GPFS customer support to request an efix, or obtain the following APAR fixes, when available:

- GPFS 2.3 - APAR IZ00619
- GPFS 3.1 - APAR IY99258

Phantom SNI devices may appear after upgrade from pre-Service Pack 7

Users Affected:

Users upgrading systems from SP6 or earlier with existing SNI devices.

Problem Description:

Changes in the HPS switch microcode and driver demand that it is absolutely necessary to install the HPS fileset updates and the GFW firmware updates without rebooting logical partitions before the system firmware is successfully installed. Rebooting the LPAR(s) prematurely will cause "phantom" SNI devices to be created on the partition(s).

The failure signature is:

1. There are twice the number of snX and sniX devices as expected in the "lsdev -C|grep sn" output
where

X = sni or sn interface number on your system anywhere from 0 thru 7 -- e.g. sni0 or sn0

2. All the sn interfaces are in the Defined state.
3. Upper half of the sniX devices are in "Available" state (higher numbered devices) and lower half of the sniX devices are in "Defined" state (lower numbered devices).
4. All the sniX devices could also be in the "Defined" state.

Common causes of "phantom" sniX devices:

- If you re-boot LPARs after upgrading HPS filesets, but prior to a successful GFW upgrade.
- Failure during GFW upgrade after updating HPS filesets.

These events/scenarios will cause all LPARs on a CEC to reboot without upgrading the GFW via AIX command line method (i.e. update_flash command).

- If you have a corrupted/truncated [GFW image file](#).
- If either /var or /tmp is too full.
- If an LPAR other than the Set Service Authority LPAR is in "Running" state during update_flash.
- If an LPAR other than the Set Service Authority LPAR is used to run the update_flash command.

Note: There may be other factors that cause LPARs to reboot after the HPS fileset is upgraded and before the GFW is successfully updated.

Recovery Procedure:

1. Complete the GFW update, and then verify that the firmware updates on all CEC(s) was successful.
2. Recover the sniX and snX devices. (Note: To recover requires at least 1 reboot. Two reboots are required if the sniX devices are busy.)
 - a. Remove all the SNI devices after making note of any customization to the SNI devices such as num_windows, driver_debug, etc . Use the "lsattr -El sniX" command so that any customization can be re-applied after the recovery procedure.

Note: This procedure resets all values back to the defaults:

To remove each logical sniX, execute the following:

```
for X in 0 1 2 3; do rmdev -d -l sni$X; done
```

If the rmdev fails for any devices (e.g. device is busy), then unconfigure the device driver as follows. Otherwise, go to step **b**.

1. Rename the configuration method for the device as follows:

```
mv /usr/sni/aix52/cfgsni /usr/sni/aix52/cfgsni.orig
```

2. Reboot each LPAR that failed.

3. Run the rmdev loop again.
4. Restore the original configuration methods name as follows:

```
mv /usr/sni/aix52/cfgsni.orig /usr/sni/aix52/cfgsni
```

- b. Remove ONLY the phantom top half of the snX devices. The lower half are real snX devices and have the ipaddr and netmask attributes in the odm. Do not delete these real snX devices.

For the top HALF of logical snX devices, execute the following:

```
for X in 2 3; do rmdev -d -l sn$X; done
```

3. Reboot the LPAR(s) --> shutdown -Fr.
4. Restore any customization to the SNI devices (e.g. num_windows, driver_debug, etc.).

hagsglsm is not reporting the local switch membership group.

Component:

RSCT

Systems Affected:

All rsct Users at Service Pack 9

Description:

In order to complete the migration of a peer domain and update the active RSCT version to a new level, you must enter the runact command as follows:

```
runact -c IBM.PeerDomain CompleteMigration Options=0"
```

This command should be run after every RSCT release upgrade.

For a more complete discussion see: "Avoiding Domain Partitioning When Migrating From RSCT 2.2.1.x or 2.3.0.x" in *IBM Reliable Scalable Cluster Technology Administration Guide*, Chapter 3. "Creating and Administering an RSCT Peer Domain." See the section on Migration.

[Back to top](#)

Restrictions

Anchor link navigation

- [GPFS maintenance level coexistence restriction](#)
- [No switch should be powered off while the SNM software is running](#)
- [Rules for swapping cables for fault isolationg](#)
- [Improved performance is more sensitive to bad links](#)
- [HPS Cluster recommended LPAR reboot procedure](#)
- [Striping mode Restrictions](#)

V3.1 maintenance levels 10 (GPFS-3.1.0.10) thru 12 (GPFS-3.1.0.12) do not coexist with other maintenance levels.

Component:

GPFS

Systems Affected:

All GPFS cluster nodes in a multi-cluster installation.

Implications:

If some nodes are running pre-3.1.0.10 GPFS and some nodes running 3.1.0.10 or later, the token manager space was filling up because of inconsistent token state between client nodes and the TM server. GPFS daemon assert Signal 11 in unlock_shark.

All nodes in the cluster must conform to one of these maintenance level compatibility restrictions:

- All nodes must be at maintenance levels 1-9 or 13 and later (GPFS-3.1.0.1 thru GPFS-3.1.0.9 or GPFS-3.1.0.13 and later).
- All nodes must be at maintenance levels 10-12 (GPFS-3.1.0.10 - GPFS-3.1.0.12).

No switch should be powered off while the SNM software is running

Component:

SNM - Switch Network Management

Systems Affected:

High Performance Switch (HPS) users applying from pre-Service Pack 7

Implications:

- A CEC frame with a switch in it CANNOT be EPOWed.
- If a CEC has to be power cycled, power down the CEC and not the frame.
- If a CEC frame with a switch needs to be EPOWed, power down the frame, kill the SNM daemon after 5 minutes, power up the frame and restart the SNM daemon.
- If one or more switches need to be recycled, power down the switches, power them back up and then recycle the SNM daemon after 5 minutes.

Rules for swapping cables for fault isolation

Component:

SNM - Switch Network Management

Systems affected:

All HPS Users applying from pre-Service Pack 7

Description:

Follow these rules when swapping cables for fault isolation:

- Only swap SNI attached cables at the switch ports to which they are attached.
- Do not swap switch to switch cables with other switch to switch cables nor with SNI attached cables.
- Do not swap cables on the SNI ports.
- Do not swap cables between switches.
- If adapters are accidentally "miswired" during the process of swapping cables, recable the adapters to their original positions.

Improved performance is more sensitive to bad links

Component:

HPS/LAPI

Problem Description:

The protocol (MPI or LAPI) will timeout if the job runs on bad links and the link routes are not fixed. If the link failure turns into adapter failure then the job gets terminated.

Solution:

To resolve this issue, monitor Service Focal Point for bad links and fix them.

HPS Cluster recommended LPAR reboot procedure

Component:

HPS/SNI

Systems Affected:

All HPS Users applying from pre-Service Pack 7

Description:

To ensure the HPS switch links are properly shutdown and re-enabled, IBM recommends that you use the following commands to recycle and reboot all LPAR/AIX images in your cluster.

1. To shut down LPAR:

shutdown -F

2. To reboot LPAR or multiple LPARs simultaneously:

shutdown -Fr

OR

dsh -av shutdown -Fr

If this procedure is not followed, then en masse reboot will certainly result in one IPC1:37:CC MP Fatal event per link.

Use of the "reboot" command or "rpower" commands will not shutdown the HPS switch links in an orderly fashion - when more than one frame at a time is cycled concurrently the SNM daemon may hang and Service Focal Point could end up with artificial errors. If when these commands must be run concurrently on multiple LPARS, it is recommended that you use them one frame at a time in your cluster.

When the HMC GUI is used to cycle an LPAR, it is recommended that you use the "shutdown" option to recycle the LPAR, not the "reset" option. The "shutdown" option will ensure that the HPS switch links

are shutdown and re-enabled cleanly - whereas the "reset" option(s) will essentially use the rpower command (and not do an orderly shutdown).

Note: Use of the rpower or HMC GUI reset options should be reserved as a "last resort" for LPARS that are not responding to a shutdown command.

Striping mode Restrictions

Component:

HPS/SNI

Description:

- Striping is currently limited to **eight** links per task. Multiple tasks can use all links as long as there are more tasks than half the links available.
- Performance degradation occurs for single task per link of 3% for unidirectional and 5% for exchange bandwidth at large messages. However, multiple tasks per link get full link bandwidth.
- In non-striping mode you only get half the links on a p655 system using a 2 plane configuration when setting MP_EUDEVICE=sn_single. You must use MP_EUDEVICE=csss or sn_all to get all links.

[Back to top](#)

Health check/planning help

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

- [Cluster support](#)
- [HPS](#)

IBM HPS POWER 4 Service Pack 22 Addendum 1 FLASH/Readme

Fix List and Readme Archive

Anchor link navigation

- [Introduction](#)
- [Supporting documentation](#)
- [Software updates](#)
- [Component updates/download information](#)

- [Installation sequence](#)
- [Installation guidelines](#)
- [Known problems and workarounds](#)
- [Fix list and readme archive](#)

HPS Service Pack fix list (by component)

The following overview lists the recommended installation sequence for installing software, firmware and operating system updates.

Anchor link navigation

- [HMC](#)
- [SNM](#)
- [GFW](#)
- [VSD](#)
- [LAPI](#)
- [HPS/SNI](#)

- [Parallel Environment](#)
- [LoadLeveler](#)

- [GPFS](#)
- [CSM](#)
- [RSCT](#)
- [Readme archive](#)

HMC

Refer to the HMC web site for specific fix information. For any fix or updated listed on a **Download** page for an HMC release, click the **View** link to see PTF-specific fix or enhancement information in the Readme for that HMC Release.

The HMC fix support website is located at the following URL:

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

SNM

For a list of abstracts see the SNM readme at the following location on the HMC web site:

<http://www14.software.ibm.com/webapp/set2/sas/f/hmc/power4/download/v33.Otherfixes.html>

Click the **View** links on this page to view service pack READMEs.

GFW

Refer to the following links for complete list of abstracts for GFW fixes in HPS Service Pack 7 or higher:

<http://techsupport.services.ibm.com/server/mdownload/7040681F.html>- for Regatta H

<http://techsupport.services.ibm.com/server/mdownload/7039651F.html>- for Regatta IH

VSD

4.1.0 fix abstracts

4.1.x.x all versions:

- System Crash in CpyRd()
- vsd.snap needs additional info

4.1.0.23:

- RVSD hangs during low paging space
- Enhance debug/data integrity code
- removevsd -fa failed to cleanup vgs on VSD primary server
- defvsd prints out extraneous info

4.1.0.22:

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

LAPI

2.4.x fix abstracts

2.4.4.1:

- Tune P6 copy routine
- IB3: lapi_init RC605 and seg fault on sles IB cluster
- Linux: PE 431 & LAPI 244 Fix Level update

2.4.3.4:

- dgsm does not handle 0 reps LAPI_DGSM_GOSUB command
- INCREASE DEFAULT TIMEOUT FOR LAPI INIT
- Linux: bump up PE & LAPI RPM fix level
- Aix Partition went to KDB

2.4.3.3:

- Change LAPI_ERR_INTERNAL5(876) LAPI_ERR_REG_TIMER(414)
- Linux: Up FIX level in LAPI_PKG_DATA to 3

2.4.3.2:

- Linux: update fix level in PE_PKG_DATA and LAPI_PKG_DATA files
- LAPI lock problem during preemption

2.4.3.1:

- Linux:Send/rev data compare failed due to LAPI_Getcptr timedout
- Jobs hang in lapi_shm_init on SLES 9
- LAPI Interrupt performance poor
- Typo for variable HAL_TS_L in lapif.h
- Linux: update PKD_DATA for PE/LAPI to show correct VRMF
- Fix LW pool init to use 'long' local variable.

2.3.3 fix abstracts

2.3.3.8:

Excessive LAPI Timer pops

2.3.3.7:

- RDMA lgpg BW ~6% lower on AIX 53E
- 2EB: mpi/us and lapi/ip job fails lapi rc=265155268

2.3.2 fix abstracts

2.3.2.8:

Excessive LAPI Timer pops

2.3.2.7:

- RDMA lgpg BW ~6% lower on AIX 53E
- 2EB: mpi/us and lapi/ip job fails lapi rc=265155268

[Back to top](#)

HPS/SNI (Devices)

1.2.0 fix abstracts

1.2.0.x all versions:

- CHGSNI caused several nodes to crash.Others nonresponsive.

- Degraded performance mode after reboot
- PATH table reader non-shipping tool.
- Performance degraded when 64k large pages are used.

1.2.0.7:

- Make the verification_test more fault tolerant
- Remove idle forked server processes
- PTF builds missing new VERIFICATION_TEST readme

1.2.0.6:

- P5 nodes crash when SNI_DD runs out of large pages
- Remove ifdef around open retry loop
- errpt msg 284BDA76 incorrectly logged

1.1.3 fix abstracts

1.1.3.x all versions:

- CHGSNI caused several nodes to crash .Others nonresponsive
- Degraded performance mode after reboot
- Performance degraded when 64k large pages are used.

1.1.3.11:

- Update automated test case output
- Make the verification_test more fault tolerant.
- Remove idle forked server processes
- PTF builds missing new VERIFICATION_TEST readme

1.1.3.10:

- P5 nodes crash when SNI_DD runs out of large pages
- Remove ifdef around open retry loop
- errpt msg 284BDA76 incorrectly logged

1.1.3.9:

- RDMA lgpg BW ~6% lower on AIX 53E
- update the hom_to_laddr map

- SNI report errpt when insufficient large pages available
- cfg error flow crash w/xmalloc_debug

[Back to top](#)

Parallel Environment

4.3.x fix abstracts

4.3.1.2

- Ckpt jobs hang in mpi_ckpt_thread_handler with PPE rfal release

4.3.1.1

- SMT processor core affinity support phase 1
- SMT processor core affinity support phase 2
- MPI_TYPE_CREATE_RESIZED in 64 bits can SEGFAULT
- PE man pages contain garbled AIX & Linux paths
- Incorrect msg matching for short MPI_Irecv from MPI_ANY_SOURCE
- PMB benchmark problem with MPI
- Incorrect msg matching for short MPI_Irecv from MPI_ANY_SOURCE
- PMB benchmark problem with MPI
- Clean up MP_TIMEOUT related code
- Linux: PE 431 & LAPI 244 Fix Level update
- Update README for PE 4.3.1 GA

4.3.0.4

- Fix for -tlp required change
- Onesided MPI tests failing in MPI_Win_unlock
- Linux: bump up PE & LAPI RPM fix level
- checkpoint/restart gives wrong result intermittently

4.3.0.3

- PMD does not recognize LoadL 3.4 for affinity
- Occasional race condition error in value returned by MPI_Iprobe
- Linux: bump up PE package FIX level

4.3.0.2

- Change default for SMT keyword
- Premature pulse timeouts with redirected STDIN and CHECKPOINT
- File /tmp/s.pedb. < poe_pid > . < uid > has permission 777
- LINUX: update fix level in PE_PKG_DATA and LAPI_PKG_DATA files
- MPCI lock leak problem

4.3.0.1

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

4.2.2.x fix abstracts

4.2.2.11:

- MPI_TYPE_CREATE_RESIZED IN 64 BITS CAN SEGFAULT
- Incorrect msg matching for short MPI_Irecv from MPI_ANY_SOURCE

4.2.2.10:

Checkpoint/restart gives wrong result intermittently.

4.2.2.7:

pct hangs at the end of the execution

4.2.2.6:

- Documentation needed on DCPLS affect on 64-bit Fortran code
- Incorrect message size returned from MPI_PROBE
- Memory leak in MPI component
- The storage used for derived data type is not freed

4.1.1.x all versions:

Loadleveler API program may get stale information when admin rETROFIT FIX FOR PMD cannot access group info

[Back to top](#)

LoadLeveler

3.4.1.x fix abstracts:

3.4.1.1

- SMT processor core affinity support phase 2
- LL linux inherits environment from RSHD/SSHD when starting PMD
- Job stuck at P state when machine names are in mixed case
- LL_GET_DATA FOR THE LL_STEPWALLCLOCKUSED, always returns 0 TIM
- IB4: LoadLever fail during initialization with IB adapters
- llchres contractions in -H help text
- llchres -F fopen only reads first argument in command list
- General parsing of integer keywords should not allow negatives
- job submission issues
- WebUI: ship with console=false
- XLOADL coredumped with customer defined CLUSTERCPUS resource
- WebUI: add to tabbed panel job builder
- Periodic Lapi Init failures on c275 cluster causing test
- fix job categories, remove health monitoring states
- RSCT 2.4.6.2 REQUIRES LOADL TO MAKE CHANGES FOR RMCAPI CALLS.
- WebUI: Total number shows wrong number in job status panel
- WebUI: widget can be enveloped by enlarging size of characters
- LOADL_NEGOTIATOR CORE DUMP FROM SERIALIZATION PROBLEM
- SMT processor core affinity support phase 2
- unexpected adapter count
- Starter fails connect to unix socket when startd runs as root
- Submit a JCF without read permission executable cause core dump
- updates for mach graphics view to complete functionality inPTF1
- Loadleveler is generating a bad dynamic block.

- mcm_affinity differences between ven and ven2
- Job not scheduling when ConsumableCpus available
- fix some bugs and improvements of job graphics view in PTF1
- FAIRSHARE PROB WITH LARGE VALUES FOR TOTAL_SHARES & INTERVAL
- error info when "initialdir" without R permission to directory
- ptf1:mcm_accumulate mcm_sni_req jobs not always schedul
- Level Bump Defect for sprelven2/ptf1
- WebUI: incomplete data in ResourcesElement
- Man page updates
- WebUI:Order of users is affected by the latest JobID
- WebUI: Number of jobs is missing when the space is not enough
- when 1 cpus_per_core used the other cpu not available
- WebUI:Info in Job Status Table is cut off
- Two blanks with class drop-down List
- WebUI: inaccurate expression of "Select Action"
- WebUI: No warning message for command that not allowed
- WebUI: Button has no effect when page is enlarged in Firefox
- WebUI: Unclear description of protocol choice
- WebUI: Button "SAVE" may remove original JCF
- preemption rset check not considering cpus that will be available
- WebUI: unknown keyword "start_date"
- WebUI: Format of "Start date" is unexpected
- No error msg as invalid date for LL version with preferences
- Hard/Soft wall clock limit's value can't be saved and submit
- WebUI: changes for jcf as script support
- Update message ELIB488 with info on setuid programs
- data can't be submitted and saved in Simple Parallel Job window
- Restore failed with Job type in Simple Parallel window
- WebUI: Help msg needed in Job Status Panel
- WebUI: Help msg needed in Machine Status Panel
- SCHEDULING PROBLEMS IN A BLUEGENE ENVIRONMENT
- LL startd crashes while jobs running
- ptf1:default mcm_affinity_options and ConsumableCpus in llstatu
- WebUI: fix fields not cleared or restored
- Update lowest AIX release needed to run LL
- WebUI: Widgets of Filters are useless
- Find PAM path failed on 64-bit platform
- invalid default limit value on starter
- experienced seven startd coredumps during NIC run

3.4.0.x fix abstracts:

3.4.0.5:

- INTERNAL AUTHENTICATION PAIR
- NEGOTIATOR CRASHING WHEN AFS PACKAGE COMMAND RUNNING
- llctl -g version gets multiple attention msg w/missing locale
- IB3 llq -l -x results in Memory Fault core dump
- Level Bump Defect for sprelven/ptf5
- LL ckpt 64bit job fails on both AIX52S and AIX53J (0709A)
- WALL CLOCK LIMIT CORRUPTED AFTER CM RECYCLE
- NEGOTIATOR HUNG
- LL: change the file system permission of socket files to 0700
- Starter fails connect to unix socket when startd runs as root

3.4.0.4:

- NEG CORE DUMPS DUE TO UNSERIALIZED ACCESS TO LLCLUSTER::NETWORK
- No support 4 Interactive POE on Linux Submit-only
- LLQ -S CORE DUMP DURING MEMORY CLEANUP
- Problem preempting 2 coschedule jobs when 1 step complete
- README Updates for PTF3
- Save option on IE does not save network/class/not
- LOADLEVELER NOT CONSIDERING THE MULTI-CLUSTER ENVIRONMENT WHEN
- preempted job resumes when not enough resource
- LLSUBMIT CHECKING EXECUTE BIT WHEN STAGING EXECUTABLE FOR CKPT
- INVALID LOADLEVELER CONFIG FILE CAUSED LL COMMANDS TO CORE DUMP
- changes to web ui for percs demo
- Remove LoadL-so-SLES9-PPC from LoadL-full-lib-ppc-*.spec file
- Venus Service Build breaks on the changes to src/ll/lib/api/wra
- Wrong VRMF for Linux Images

3.4.0.3:

- LOADLEVELER CONSTRAINS CONSUMABLE RESOURCES WHICH ARE NOT BEING
- JCF KEYWORD VALUES CONTAINING LL VARIABLES CANNOT BE OVERRIDDEN
- Memory leaks caused by callers of itoa
- MEMORY LEAK IN LL GET DATA API WHEN AFFINITY IS CONFIGURED
- A job in R state but actually stuck in Vacate Pending for 2 hrs
- MEMORY ERROR ALLOCATING MEMORY TO HOLD AFS TOKENS
- TOKENS CMD DOES NOT RETURN THE CORRECT UID
- LoadL PTF needs to show correct level

3.4.0.2:

- Change default for SMT keyword
- Eliminate current limitations running NIS/CHECKPOINT handler
- CM coredumped in _int_malloc (malloc.c) while testing IB
- mpichgm various problems with mpirun options
- 4net Using instances=8, LL only displays 2 of 4 planes
- LOADL jobs at same time can result in non-unique temp name
- Jobs stuck in ST state on the C559 system
- User prolog error msg when failure to setup user environment
- LoadLeveler API must be linked with XLC
- StartD status for Draining shows Unknown
- Job stuck in VP state
- If bad nodes/tasks and save, does not tell you not-saved
- Deferred state should be moved to Running-like
- Refresh action on machines updates timestamp for jobs as well
- Console below submit dialog should default to last entry view
- Change default for SMT keyword
- su permission issue / env var not preserved on file save
- Help text fixes/design in portlets
- CK checkpointing status not showing on Job Status window
- tool tip missing from check boxes
- LOADL negotiator receives SIG 4 when starting job with blockin
- StartD/SchedD_runs_here values returned for UI display
- File system full says OEB
- MPI jobs (started with poe) got strange msg in stderr & stdout
- misc webui fixes, multi job/machine selection capability needed
- LOADL 3.3.2.5 data access API can core dump with 64-bit apps
- Memory leak/LOADL 3.3.2.5 data access API linking 32-bit apps
- Startd ignores VACATE order if does not have step object
- llstatus -v fails if config files have error
- IT Hierarchical job start error unknown with hwlmmn
- job selection missing from webui graphic view
- create WLM class for suspended ConsumableMemory steps
- /usr/lib/libc.a(shr.0) AIX only
- fix webui graphic job selection
- Message catalog cleanup
- BG: LL removes pre-defined block if reuse is off & preemption
- LL can allocate a BG partition to a job not de-allocated yet
- fix job running hilite
- Job stuck in RP state when remove cmd follows complete cmd
- Job stay in ST state if start & kill cmds reach node same time
- job stuck in VP state during high stress run
- LoadL PTF needs to show correct level for ptf set 2
- update builder fields, version

- webui graphic machine missing popups
- VEN PTF2 build fails on Packaging. Packaging part files for
- Avoid routing unwanted adapter_usage attributes while Job start

3.4.0.1:

- AFS KTC Token Size Compatibility
- Problem during preemption when ifconfig adapter down
- LIMCluster leak during Schedd reconfig
- A LOADL recycle on the central manager when a reservation is
- llq -s gives erroneous message when any schedd daemon is down
- job remains in ep too long
- Missing serialization in LIMcm::updateAdapterList
- Memory leak in sch_api_ext.c sample program when executed
- Inc msg from llq -s when not enough initiators and need adapte
- LL mail not delivered due to invalid mail address
- Memory leaks in LoadL_negotiator daemon on BGL
- llmovespool does not work on RH4 system
- remove NQS entries from llq -l and llq -l -x listings
- MetaCluster JobID displayed in long listing of non_meta jobs
- Problem running a 3.3.1.3 and 3.3.2.4 mixed cluster
- builder, correct locs, fields, actions
- New LL cmd option -W should work for llclass max & free slots
- llq and llqres display incorrect header and separator
- Problem running multistep job with mpich-gm and poe job steps
- AFS RH4 logging and README errors
- adapter count(0) returned by LL
- Hold status for job displays as state D, errantly
- sample myclass.c fails compile
- coschedule jobs and preemption
- jobs not restarting correctly when stop LL and then start again
- LoadLeveler does not try reconnecting to rsct- rmc_start_sessi
- Schedd/Startd states incorrectly reported
- LoadLeveler logfile being created with other write permission
- Mpich job command files needs to changed to use new env variabl
- Sample file changes for quick start implementation
- Using llmovejob the remote filter is not called but local is
- Adapter coexistence bug submitting a job from LL 3.3.1 TO 3.3.
- Coexistence: llq can't show job info submitted from rmer
- update packaging for sles10 to include llmovespool
- Miss version check in multicluster connection
- LoadL PTF needs to show correct level for ptf set 1
- Startd core dump on IB Linux cluster.

- mvapich sample jcf missing
- path not set for commands on job queue
- Job selection is unknown to user, affects filter views
- Save filename for builder must specify full pathname
- Packaging issue with LL Web UI
- Save option from Create&Submit does not save network/class/not
- Node and task usage for Create and Save
- Error messages from save not passing when directory unwritable
- Debug needed for determining data from/to LL,JNI,etc.
- documentation-readme and welcome page
- update post_u.sh for llwebui
- Redundant data in hover-help for Job/Machine status
- Only last step is shown in multi-step jobs

3.3.2.x fix abstracts:

3.3.2.x all versions:

- Eliminate current limitations running NIS/CHECKPOINT handler
- CM coredumped in _int_malloc (malloc.c) while testing IB
- 4net Using instances=8, LL only displays 2 of 4 planes
- Load leveler API must be linked with XLC
- Loadl negotiator receives sig 4 when starting job with blockin
- LOADL 3.3.2.5 data access api can core dump with 64-bit apps
- MEMORY LEAK/LOADL 3.3.2.5 data access api linking 32-bit apps
- Create WLM class for suspended ConsumableMemory steps
- /usr/lib/libc.a(shr.0) AIX only
- LL can allocate a bg partition to a job not de-allocated yet
- Loadlever constrains consumable resources which are not being
- JCF keyword values containing ll variables cannot be override
- Update vrmf file for LoadL 3.3.2 PTF8

3.3.2.8:

- Eliminate current limitations running NIS/CHECKPOINT handler
- CM coredumped in _int_malloc (malloc.c) while testing IB
- 4net Using instances=8, LL only displays 2 of 4 planes
- LOADL API must be linked with XLC
- LOADL negotiator receives SIG 4 when starting job with blockin
- LOADL 3.3.2.5 data access API can core dump with 64 bit apps
- Memory leak/LOADL 3.3.2.5 data access API linking 32 bit apps
- create WLM class for suspended ConsumableMemory steps

- /usr/lib/libc.a(shr.0) AIX only
- LL can allocate a BG partition to a job not de-allocated yet
- LoadLeveler constrains consumable resources which are not being
- JCF keyword values containing LL variables cannot be override
- Update vrmf file for LoadL 3.3.2 PTF8

3.3.2.7:

- Problem during preemption when ifconfig adapter down
- A LOADL recycle on the central manager when a reservation is
- Memory leak in sch_api_ext.c sample program when executed
- LOADL jobs at same time can result in non-unique temp name
- User prolog error msg when failure to setup user environment
- LoadLeveler does not try reconnecting to rsct- rmc_start_sessi
- LoadLeveler logfile being created with other write permission
- Using llmovejob the remote filter is not called but local is
- Bad Function Call Parameter reported by ZeroFault
- Adapter coexistence bug submitting a job from LL 3.3.1 TO 3.3.
- Submit job on AIX will failed due to reservation error.
- Level Bump Defect for sprelmer/ptf7

3.3.2.6:

- llq -f and llq -r should return favored_job status as an option
- Ability to query LL version in MultiCluster env
- llclass -X broken between cs3c2 and mercury
- job remains in ep too long
- On BLUE GENE the llcancel can take a long time
- The llmodify -w does not work properly with the backfill sched
- Incorrect casting of variables in LoadL_schedd
- Linux: Wrong secondary groups used at run time
- Negotiator deadlock scheduling SP2 switch
- Inc msg from llq -s when not enough initiators and need adapte
- Creating Reservation causes Negotiator to crash
- LL mail not delivered due to invalid mail address
- Problem running a 3.3.1.3 and 3.3.2.4 mixed cluster.
- llq and llqres display incorrect header and separator
- LoadL PTF needs to show correct level for ptf set

3.2.0.22:

- Multiple updates received at once causes incorrect adapter DAT

[Back to top](#)

GPFS

Refer to the following change logs:

3.1.x

- [GPFS 3.1.13](#)
- [GPFS 3.1.12](#)
- [GPFS 3.1.11](#)

2.3.x

- [GPFS 2.3.22](#)
- [GPFS 2.3.21](#)
- [GPFS 2.3.20](#)

CSM

1.5.1.x fix abstracts

1.5.1.6:

- Another process can grep RMCS port (657) before CSMFIRSTBOOT R
- AIX: No CSM post install activities and install.log
- RH:dsh failed because known_hosts didn't update after install
- ANother process can grab RMCS port (657) before OSFIRSTBOOT
- CT: MM peer domain partitioned and can not merge back
- Varying type byte in Crichton VPD causes lshwstat uuid to fail
- CSM DSH command runs out of file descriptors
- bmc2 lib causing HwCtrlRM to core dump
- ERRMSGQUE overflow can cause node to hang
- dsh hangs when we run updatenode -a with more than 2 nodes
- Update VRMF and READMEs for CSM 1.5.1.6

1.5.1.5:

- updatehwdev failed when running on a pLinux MS to a HMC

- RPOWER TO OPEN_FIRMWARE prompt fails in full partition mode
- DSH -S option truncates output characters
- UPDATENODE -C doesnt configure etherchannel adapters
- DPING on nodes broken
- SOL remote console fails for e326m
- Update VRMF and READMEs for CSM 1.5.1.5

1.5.1.4:

- CFMUPDATENODE creates a dir with bad perms if ._hostname exist
- Fix dsh/dcp SMIT panels
- RFWSCAN should unset DSH_PATH
- getserviceableevents should delete DSH_PATH
- DMSRM core dump during updatenode
- New BMC lflash package for x346 fails with rfwflash
- rconsole xterm window size wrong if -o option used
- DSH no longer handling multiple commands from a standard input
- Update VRMF and READMEs for CSM 1.5.1.4
- INSTALLNODE runs cfm on all the nodes after first boot

CSM 1.4.1.15:

- DSH no longer passes -N TO RSH which breaks the CSMSETUPNIM command checks for hostnames
- IBM.HWCTRLRM daemon has memory leakage in a CSMAgentRM to refetch device tree dump during fld:CSMAgentRM ran out memory after being
- DSH broken with redirected standard input
- Update VRMF for csm_rcat PTF 15
- LX: IBM.DMSRM core dump during application of ssh setup failure

[Back to top](#)

RSCT

2.4.6.x fix abstracts

2.4.6.3:

- RSCT all NIMS die at once lead to partion HACMP cluster
- File named used for monitoring may cause problem
- Sometimes GblRM fails to startup during node online

- RMCD CORE WHEN RUNNING PORT SCANNER
- Domain goes offline if TOD goes backwards during propagation
- Differing DeviceName not respected
- deadlock in destructRcp inside enumeratRcps
- ORACLE INSTANCE CRASH - ISSUE WITH HAGSD SUBSCRIBE/UNSUBSCRIBE
- null pointer passed to lockInt()
- LX:duplicate exit() call cause core
- Manual unmount of a AgFileSystem results in OpState 3
- CONFIGRM CG COMBINING IS AFFECTED BY MULTIPLE ROUTES
- rmcd disconnect RM due to alloc fail in read_stream_soc
- HATS_DISKHB_NIM CORE - 256MB ADDRESSING SPACE FOR NIM EXHAUSTED
- TSA: CritRsrcActive takes to long to transit from 1 to 0
- Set VRMF for Dvorak PTF 3
- CONFIGRM CG COMBINING IS AFFECTED BY MULTIPLE ROUTES
- Mutiple trace file merge produce the garbled-order output
- /VAR/CT/IW/SOC/CTCAS/CLIENT HAS A GROWING NUMBER OF SOCKET FIL
- RSCT.OPT.STORAGERM FAILS INSTALL IN MULTIBOS ENVIRONMENT
- CritRsrcActive equals 1 in falsely after SRM daemon restarted
- ConfigRM cores when startdomain on 100 nodes
- Core on non-LVM Linux systems due to empty command string

2.4.6.2:

- PeerNode:OpState: no event after node come back online
- ConfigRM behavior when mkrpdomain fail due to /var full
- MM: do not add the invalid address to THL
- GlobalRM: need graceful handle when GS down, no core please
- correct some memory leak leftover
- CT: cthagsglsm core dump after apply DVO ptf on COP
- CDL: GFVTRM randomly core at exit
- Resources AgFileSystem are not monitored as required for SA
- GlobalRM core if rmcd die while query IBM.PeerNode
- hags core at SVProtocal.C due to access address 0x0
- CDL: GFVTRM abort at CRunnable
- CT:LX: trace files are unnecessarily reset when node number cha
- AIX multiple disks per VG 7
- Notify RMC daemon when /var/ct/cfg/clusters file changes.
- lsrsrchrsrc help msg no longer shows -c -a as a valid combo
- Set VRMF for Dvorak/Dvorak52 PTF 2
- CT: CRunnable::threadMain() not catching all exceptions
- Remove hooks to detect cluster file changes from the RMC daemon
- IBM.HOSTRM:MKCONDITION CANT USE BOTH RUSER &
- add iwaccess to Boson/Quark not to break regression

- Trace exit handler doesn't work with dlclose
- unable to online FS & VG's State is STUCKED
- allow rmc to be down when notifying it of .../clusters change.
- CDL: One redundant pthread_mutex_unlock call.
- Unable to online exported vg
- reset of AgFS leaves VG rsrc in Stuck Online state
- CDL: for StorageRM Diagnostic Tool checking in
- Agg rsrc created in 1 node domain when rsrc becomes a ghost
- CDL: Kill GS cause glsm hang up for 65~70 sec.
- can online a resource when odm doesn't match disk info
- GS_ERROR_ER AFTER APPLYING RSCT 2.4.6 OR 2.3.10
- FS Name is different if Label is changed
- Unable to create IBM.AgFileSystem resource
- Timing issue during automated move of file system
- pclose missing
- track138458 - AutoMonitor doesn't work for LVM agfs
- Deadlock on "chsrc IBM.AgFileSystem AutoMonitor=1"
- Linux AgFS AutoMonitor rsrc onlines with any manual mount point
- startdomain with man.mounted fs causes RM

2.4.6.1:

- keep shm size as reasonable limit in large pagesize
- User able to change const Name in harvested IBM.AgFS
- IBM.CONFIGRM USING TOO MUCH CPU DUE TO HAGSGLSM MSG QUEUE ISSU
- nfs filesystem flipflops between online and offline when startd
- CsiAttrs shouldn't be marked as a req attr for SAF_CSI
- LX: Able to change the Name attr of the IBM.VolumeGroup rsrc
- possible memory leak in RMACCommandGroup.C
- HAGSD MIGHT ASSERT WHEN RECEIVING NON-RSCT PACKETS
- forceroffline should fail if IW is specified as an operand
- CDL: StorageRM core found on c108f1n15.
- cthags coredumps on Zorro Linux cluster in prm code
- startdms does not work
- problem with softdog and critical TSA resources on SLES10
- CtSec segfault on corrupted THL file
- CDL: chsrc-api core dumps when invoked with -I
- hags core dump due to unnamed block in TraceStream.operator
- CT: rmc core dump abort core dump while apply dvo to debs003a
- Do not ship symlink in /usr/lib/ to libct_amf.so
- AGFileSystem not failed offline when pulling cable
- Resources AgFileSystem on same VG only partially online
- RM might die when N-N batch has fewer targets than commands

- More cthags coredump after scaling RPD on Zorro
- Initial OpState for the stopped node should be OFFLINE
- ppc64 RHEL5 Beta1: rsct daemons can not be started
- CDL: chsensor returns an incorrect hint.
- CDL: startsrc doesn't fail when mount dir doesn't exist.
- LX: call cu_pick_thread_stacksize() when creating security thre
- Unable to change ProtectionMode or DeviceLockMode in VG or LV
- Bump VRMF for Dvorak/Dvorak52 PTF 1
- attach shm fail in large page sz
- call cu_pick_thread_stacksize for proper adjustment
- CT: ctmskf core dump at clic.CLiC_dispose
- CRunble: adjust thread stack size to reasonable amount
- FIX TO PHOENIX.SNAP CODE UNDER SNAP -E.
- Incorrect SysMountPoint is harvested
- "Dead Man Switch Disabled" displays on console when node online
- c559f11sq01: hang in ct2prm functions
- CT: deadlock in RMRccp::enumerateRcps()
- Correct the CMVC Prolog in SAFRM mdef file

2.3.11.x fix abstracts

2.3.11.1:

- lti-node DHB networks: DHB NIM + dhb_reader
- rmcd cored at pci_prm_first_init
- NODE FAILS TO REJOIN RPD DUE TO GLOBAL VERSION MISMATCH
- lsrpdomain command hangs
- output not the same
- GL ConfigRM core at PeerDomainRcp::selectCSSK
- ctsnap doesn't get lssrc -ls output on AIX
- selectCSSK handles unreachable nodes badly
- ConfigRM hit core in scalar test.
- WLMRM deadlock at RMRccp::unbindRCP
- ctrl operation support
- AIX: SDD driver disk monitoring fails
- StorageRM core file due to buffer overflow
- HATS hit core during IB regression.
- the CSSKTYPE_DES_MD5 version is incorrect
- cthats cored at RPDSecStatus
- null out both length and value members of sec_buffer_desc
- hagsglsmid hit core during IB regression.
- bad indexing in SelCSSKEncrCache.C
- Set VRMF for Elgar/Elgar52 PTF 1

- interop: Domain partitioned after UpdateKey
- interop: lsrc -ls IBM.ConfigRM shows security disabled
- rsct.basic.rte msgs during elg52 apply
- nodes repeatly online / pending online
- correct error handling on failed selectCSSKOps
- domain pending online if mkrpdomain with -k CSSKTYPE_None
- interop: startprnode failed to bring all nodes online
- fix forcerpoffline error msg EMsgforcerpofflineRMCErr
- RMCD CORES WHEN INTERNET SCANNER 7.0 (ISSCAN) USED
- Several messages with invalid format in IBM.StorageRM.msg
- CSSKRefreshInterval should = 0 when cssk changes to None
- caGetPMFile fails on partition
- IBM.FSRMD CORE DUMP DUE TO MEMORY LEAK
- core passing NULL pointer
- DeviceLockMode for MD rsrc changed from 0 to 4
- interop: CSSK not refreshed periodically
- Client Monitoring Interval set to 0 when no CGs are defined
- old level of chkconfig affects start order of RSCT daemons
- LX: RSCT Elgar man page updates
- CONFIGRM CORE DUMP DUE TO CSSKTYPE IS NOT SET.
- DeviceLockMode fixes for AIX/Linux
- CSSKRefreshInterval is not set to 0 after ptf apply

2.3.10.x fix abstracts

2.3.10.3:

- RSCT all nims die at once lead to partion HACMP cluster
- File named used for monitoring may cause problem
- Sometimes GblRM fails to startup during node online
- RMCD core when running port scanner
- Domain goes offline if TOD goes backwards during propagation
- Differing DeviceName not respected
- deadlock in destructRcp inside enumeratateRcps
- ORACLE instance crash - issue with HAGSD SUBSCRIBE/UNSUBSCRIBE
- null pointer passed to lockInt()
- LX:duplicate exit() call cause core
- Manual unmount of a AgFileSystem results in OpState 3
- CONFIGRM CG combining is affected by multiple routes
- rmcd disconnect RM due to alloc fail in read_stream_soc
- HATS_DISKHB_NIM CORE - 256MB ADDRESSING SPACE FOR NIM EXHAUSTED
- TSA: CritRsrcActive takes to long to transit from 1 to 0
- Set VRMF for Dvorak PTF 3

- CONFIGRM CG combining is affected by multiple routes
- Mutiple trace file merge produce the garbled-order output
- /VAR/CT/IW/SOC/CTCAS/CLIENT HAS A GROWING NUMBER OF SOCKET FIL
- RSCT.OPT.STORAGERM fails install in MULTIBOS environment
- CritRsrcActive equals 1 in falsely after SRM daemon restarted
- ConfigRM cores when startdomain on 100 nodes
- Core on non-LVM Linux systems due to empty command string

2.3.10.2:

- PeerNode:OpState: no event after node come back online
- ConfigRM behavior when mkrpdomain fail due to /var full
- MM: do not add the invalid address to THL
- GlobalRM: need graceful handle when GS down, no core please
- correct some memory leak leftover
- CT: cthagsglsm core dump after apply DVO ptf on COP
- CDL: GFVTRM randomly core at exit
- Resources AgFileSystem are not monitored as required for SA
- GlobalRM core if rmcd die while query IBM.PeerNode
- hags core at SVProtocol.C due to access address 0x0
- CDL: GFVTRM abort at CRunnable
- CT:LX: trace files are unnecessarily reset when node number cha
- AIX multiple disks per VG
- Notify RMC daemon when /var/ct/cfg/clusters file changes.
- lsrsrc/chrsrc help msg no longer shows -c -a as a valid combo
- Set VRMF for Dvorak/Dvorak52 PTF 2
- ctcasd still need LD_ASSUME_KERNEL=2.4.19 for RHEL3
- CT: CRunnable::threadMain() not catching all exceptions
- Remove hooks to detect cluster file changes from the RMC daemon
- IBM.HOSTRM:MKCONDITION CANT USE BOTH RUSER &
- add iwaccess to Boson/Quark not to break regression
- Trace exit handler doesn't work with dlclose
- unable to online FS & VG's State is STUCKED
- allow rmcd to be down when notifying it of ../clusters change.
- CDL: One redundant pthread_mutex_unlock call.
- Unable to online exported vg
- reset of AgFS leaves VG rsrc in Stuck Online state
- CDL: for StorageRM Diagnostic Tool checking in
- Agg rsrc created in 1 node domain when rsrc becomes a ghost
- CDL: Kill GS cause glsm hang up for 65~70 sec.
- can online a resource when odm doesn't match disk info
- GS_ERROR_ER AFTER APPLYING RSCT 2.4.6 OR 2.3.10
- FS Name is different if Label is changed

- Unable to create IBM.AgFileSystem resource
- Timing issue during automated move of file system
- pclose missing
- track138458 - AutoMonitor doesn't work for LVM agfs
- Deadlock on "chsrc IBM.AgFileSystem AutoMonitor=1"
- Linux AgFS AutoMonitor rsrc onlines with any manual mount point
- startdomain with man.mounted fs causes RM

2.3.10.1:

- keep shm size as reasonable limit in large pagesize
- User able to change const Name in harvested IBM.AgFS
- IBM.CONFIGRM USING TOO MUCH CPU DUE TO HAGSGLSM MSG QUEUE ISSU
- nfs filesystem flipflops between online and offline when startd
- CsiAttrs shouldn't be marked as a req attr for SAF_CSI
- LX: Able to change the Name attr of the IBM.VolumeGroup rsrc
- possible memory leak in RMACCommandGroup.C
- HAGSD MIGHT ASSERT WHEN RECEIVING NON-RSCT PACKETS
- forceroffline should fail if IW is specified as an operand
- CDL: StorageRM core found on c108f1n15.
- cthags coredumps on Zorro Linux cluster in prm code
- startdms does not work
- problem with softdog and critical TSA resources on SLES10
- CtSec segfault on corrupted THL file
- Bump hmc.lparcmdrm to 1.3 for eClipze
- CDL: chsrc-api core dumps when invoked with -I
- hags core dump due to unnamed block in TraceStream.operator
- CT: rmc core dump abort core dump while apply dvo to debs003a
- Do not ship symlink in /usr/lib/ to libct_amf.so
- AGFileSystem not failed offline when pulling cable
- Resources AgFileSystem on same VG only partially online
- RM might die when N-N batch has fewer targets than commands
- More cthags coredump after scaling RPD on Zorro
- Initial OpState for the stopped node should be OFFLINE
- ppc64 RHEL5 Beta1: rsct daemons can not be started
- CDL: chsensor returns an incorrect hint.
- CDL: startsrc doesn't fail when mount dir doesn't exist.
- LX: call cu_pick_thread_stacksize() when creating security thre
- Unable to change ProtectionMode or DeviceLockMode in VG or LV
- Bump VRMF for Dvorak/Dvorak52 PTF 1
- attach shm fail in large page sz
- don't set LD_ASSUME_KERNEL on RHEL5
- call cu_pick_thread_stacksize for proper adjustment

- CT: ctmskf core dump at clic.CLiC_dispose
- CRunble: adjust thread stack size to reasonable amount
- FIX TO PHOENIX.SNAP CODE UNDER SNAP -E.
- Incorrect SysMountPoint is harvested
- "Dead Man Switch Disabled" displays on console when node online
- hang in ct2prm functions
- CT: deadlock in RMRccp::enumerateRcps()
- Correct the CMVC Prolog in SAFRM mdef file.

2.3.10.x fix abstracts

2.3.7.1:

- lti-node DHB networks: DHB NIM + dhb_reader
- rmcd cored at pci_prm_first_init
- NODE FAILS TO REJOIN RPD DUE TO GLOBAL VERSION MISMATCH
- lsrpdomain command hangs
- output not the same
- GL ConfigRM core at PeerDomainRcp::selectCSSK
- ctsnap doesn't get lssrc -ls output on AIX
- selectCSSK handles unreachable nodes badly
- ConfigRM hit core in large scalar test.
- WLMRM deadlock at RMRccp::unbindRCP
- MD: ctrl operation support
- AIX: SDD driver disk monitoring fails
- StorageRM core file due to buffer overflow
- HATS hit core during IB regression
- The CSSKTYPE_DES_MD5 version is incorrect
- cthats cored at RPDSecStatus
- null out both length and value members of sec_buffer_desc
- hagsglsm hit core during IB regression.
- bad indexing in SelCSSKEncrCache.C
- Set VRMF for Elgar/Elgar52 PTF 1
- interop: Domain partitioned after UpdateKey
- interop: lssrc -ls IBM.ConfigRM shows security disabled
- rsct.basic.rte msgs during elg52 apply
- nodes repeatly online / pending online
- correct error handling on failed selectCSSKOps
- domain pending online if mkrpdomain with -k CSSKTYPE_None
- interop: startprnode failed to bring all nodes online
- fix forcerpoffline error msg EMsgforcerpofflineRMCErr
- RMCD CORES WHEN INTERNET SCANNER 7.0 (ISSCAN) USED
- Several messages with invalid format in IBM.StorageRM.msg

- CSSKRefreshInterval should = 0 when cssk changes to None
- caGetPMFile fails on partition
- IBM.FSRMD CORE DUMP DUE TO MEMORY LEAK
- core passing NULL pointer
- DeviceLockMode for MD rsrc changed from 0 to 4
- interop: CSSK not refreshed periodically
- Client Monitoring Interval set to 0 when no CGs are defined
- old level of chkconfig affects start order of RSCT daemons
- LX: RSCT Elgar man page updates
- CONFIGRM core dump due to CSSKTYPE is not seT
- DeviceLockMode fixes for AIX/Linux
- CSSKRefreshInterval is not set to 0 after ptf apply

[Back to top](#)

HPS P4 README Archive

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

[IBM HPS POWER4 FLASH/Readme - Service Pack 22](#)

[IBM HPS POWER4 FLASH/Readme - Service Pack 21](#)

[IBM HPS POWER4 FLASH/Readme - Service Pack 20](#)

[IBM HPS POWER4 FLASH/Readme - Service Pack 19](#)

[IBM HPS POWER4 FLASH/Readme - Service Pack 18](#)

IBM HPS POWER4 FLASH/Readme - Service Pack 17 is not available

[IBM HPS POWER4 FLASH/Readme - Service Pack 16](#)

[IBM HPS POWER4 FLASH/Readme - Service Pack 15](#)

[IBM HPS POWER4 FLASH/Readme - Service Pack 14](#)

[IBM HPS POWER4 FLASH/Readme - Service Pack 13](#)

[IBM HPS POWER4 FLASH/Readme - Service Pack 12](#)

[IBM HPS POWER4 FLASH/Readme - Service Pack 11](#)

[IBM HPS POWER4 FLASH/Readme - Service Pack 10](#)

[IBM HPS POWER4 FLASH/Readme - Service Pack 9](#)

[IBM HPS POWER4 FLASH/Readme - Service Pack 8](#)

[IBM HPS POWER4 FLASH/Readme - Service Pack 7](#)

[IBM HPS POWER4 FLASH/Readme - Service Pack 6](#)

[IBM HPS POWER4 FLASH/Readme - Service Pack 5](#)

[Back to top](#)

Health check/planning help

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