This document remains available for reference purposes only. All external links have been removed. Please refer to the latest service pack document for interactive capabilities and available software.

HPS Service Pack 8 FLASH/Readme First

IBM pSeries High Performance Switch (HPS) Service Pack 8 requires updates to the Hardware Management Console (HMC), Switch Network Manager (SNM) also referred to as FNM on the HMC, Global Firmware (GFW), Power Subsystem Microcode (ptcode), AIX 5L Version 5.2 and various supporting AIX LPPs.

This document contains general guidelines for upgrading the components listed in Component update/download information. **These guidelines are intended to be a supplement to the other IBM documents referred to in this document.** We strongly advise that you have the referenced documents available before you begin the upgrade process.

Currently only the levels listed in Component update/download information are supported with HPS Service Pack 8. Therefore, if you are migrating/installing Service Pack 8, all components should be at the indicated service levels when you are finished. No mixing of Service Pack maintenance levels is supported at this time.

Therefore, if you are migrating/installing Service Pack 8, all components should be at the indicated service levels when you are finished. No mixing of Service Pack maintenance levels is supported at this time.

Users must follow the procedure in the order outlined in Section 2: "Recommended Installation Sequence (Overview)". Failure to do so may cause problems with configuring the HPS SNI devices.

Because '/var' is a system data repository, sys admins should check periodically to maintain /var such that it is under 50% full. [use 'df -k']
If it is more than 75% full, look for the directories that contain the most data. [use 'du /var | sort -n']

Table of Contents:

Section 1: Component update/download information

Component	Service pack updates	Download sites
HMC	This service pack requires that the HMC is updated to HMC V3 R3.0: --- HMC_Update_V3R3.1.zip	HMC corrective service Note: The 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. --- http://techsupport.services.ibm.com/server/hmc/power4/fixes/mcode/r3v30.html
SNM/FNM	APAR IY58991 PTF U499536 Choose "SNM_Serv_Pack8.zip"	HMC corrective service
GFW	3H040602.img (p690) or 3J040602.img (p655) Go to the "System Microcode" section. For 7040-681 Version 3 p690 models, choose "version 3H040602". For 7039-651 p655 models, choose "version 3J040602". Note that GFW is available in IBM CORE 3-4 days earlier then the above mentioned website. Please contact your IBM CE for the GFW in IBM CORE if not available on website. Refer to this website for detailed download and unpacking procedures: http://techsupport.services.ibm.com/server/mdownload/downproc.html	Microcode download (System)

Power Subsystem Microcode	ptcode-1.67.2665-1.i386.rpm (Power code is the same for both p690 and p655) Choose "Power Subsystem Microcode for 7039-651 (p655) and servers containing the 7045-SW4 (High Performance Switch)". Note that Power Subsystem Microcode is available in IBM CORE upto 1 week earlier then the above mentioned website. Please contact your IBM CE for the Power Subsystem Microcode in IBM CORE if not available on website.	Download microcode (Other)
AIX	The recommended AIX service level for AIX 5L version 5.2 is the 5200-03 Recommended Maintenance package (APAR IY56769)	pSeries support fixes

AIX LPPs:
If you have any of the following AIX LPPs Installed, apply the listed APARs, which are needed for HPS Service Pack 8. They are available from the following location: pSeries support fixes

A general way to check levels, issue:
`lslpp -l | egrep "vsd|LAPI|HPS|sni|ppe|LoadL|mmfs|rsct|csm|essl|pessl`
and then compare the results to the **Level check** column in the following table.

LPP	Release	Component ID	APAR	Level check
VSD	410	5765G2602	IY58792	rsct.vsd.cmds 4.1.0.8
LAPI	231	5765G2601	IY58795	rsct.lapi.rte 2.3.1.8
HPS/SNI	110	5765G2400	IY58796	devices.chrp.IBM.HPS.rte 1.1.1.1
PPE	410	5765F8300	IY58797	ppe.poe 4.1.0.9
LoadL	320	5765E6900	IY58798	LoadL.so 3.2.0.8
GPFS	210	5765F6400	IY58799	mmfs.gpfs.rte 2.1.0.15
GPFS	220	5765F6400	IY58800	mmfs.base.rte 2.2.1.0
CSM	133	5765E88AP	IY58801	csm.client and csm.server 1.3.3.2
RSCT	233	5765F07AP	IY58802	rsct.basic.rte 2.3.3.3
Parallel	310	5765F8400	PQ86328	pessl.rte.common 3.1.0.1
ESSL			PQ86327	pessl.rte.rs1 3.1.0.1
			PQ86327	pessl.rte.smp 3.1.0.1
ESSL	410	5765F8200	N/A	essl.rte.common 4.1.0.0 essl.rte.rs1 4.1.0.0 essl.rte.smp 4.1.0.0

Note, this is not intended to be a complete list of fileset updates in this Service Pack. For the complete list, refer to the above download website.

Section 2: Recommended installation sequence (overview):

1	Install HMC R3 V3.0 base and driver R42H-82-HMC-20040505.1 on all HMCs in the system - Required with Service Pack 8.
2	Install HMC PTF HMC_Update.V3R3.1.zip - Required with Service Pack 8. See detailed instructions in the HMC Installation Guidelines in Section 3.
3	Install SNM_Serv_Pack8.zip - Required on the HMC with Service Pack 8. The SNM contained and installed as part of the new HMC R3 V3.0 base is not the current level.
4	Reboot HMC to complete SNM software update.
5	Disable SNM software from the HMC GUI.
6	Apply HPS/SNI LPP fileset "update" images to LPARs that have the HPS 1.1.1.0 base images installed (see "AIX LPPs" in Section 1: Update/Download Information) IMPORTANT: Do not reboot logical partitions until after Step 8 is complete! Rebooting prematurely will generate "phantom" SNI devices. See Problem #1 in the "Known Problems" section of this document for more information.
7	For each CEC on which you want to install GFW code, if using the recommended AIX command line (update_flash) method, shutdown all partitions except the one with service authority.
8	Install the appropriate GFW driver on CEC(s).
9	After completion of the GFW install - enable SNM from the GUI.
10	If you installed GFW from the AIX command line - i.e. using update_flash then power down CEC(s) using CSM rpower or GUI method - not from EPO red switch.
11	Power up CEC(s) and activate partition(s).
12	Enable technical large page support - Required for levels greater than Service Pack 6. Note: Refer to "HPS/SNI" sub-section in the "Installation Instructions" section below for details on technical large page setup.
13	Power cycle switch planes that were just upgraded (required for this Service Pack release).
13a	Recycle SNM software from the HMC GUI.
14	Install AIX base updates.
15	Install the AIX LPP updates and reboot partition(s).

Section 3: Installation instructions (detailed)

HMC Important Notes:

- Install the HMC code by following the instructions in the pSeries High Performance Switch (HPS) Planning, Installation and Service Guide (GA22-7951-01), which will be referred to as the HPS guide for the rest of this document.
- Have your IBM CE download the most recent copy of the HPS guide from IBM CORE to get updated HPS install information.
- Also you should review the HMC information on the web page where you downloaded the images.
- Service Pack 8 REQUIRES HMC V3.3.0 Driver 82 Rev1 (required since Service Pack 6 release) and HMC V3.3.1 PTF update.
- HMC V3.3.0 Driver 82 Rev1 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 PTCode on your HMC and you will have to reaquire it for future installs.
- The HMC is now installed using 2 CD's - RG_BASE_040507.1.iso and R42H-82-HMC-20040505.1.iso.
- Upgrade installation is only supported when upgrading from HMC 3.2.X.
- The login available at virtual console 0 (via the CTRL-ALT-F1 key sequence) is no longer available.
- Only New Installation and Upgrade Installation are supported for this release.
- 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.
Perform a Save Upgrade Data task from the HMC console.

Note: The save upgrade data task should be run immediately before reboot of the HMC with the recovery CD. If the HMC reboots and does not go to the install menu of the recovery CD, you should repeat the save upgrade task. For more information see:
https://techsupport.services.ibm.com/server/hmc/power4/fixes/mcode/ptf_r3v30mc.html

- The HMC_Update.V3R3.1.zip must be installed once the new HMC installation or upgrade procedure is complete.

Installation GUIDELINES: (Refer to the most recent HMC Documents for complete procedure details.)
Order the Hardware Feature Code (MES) 0960 for the initial upgrade CDs from your IBM Sales Representative
or Business Partner.
As part of any system change - it is recommended to have 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

Perform a Save Upgrade Data task. This task is under the Software Maintenance folder on the HMC console.
Note: The save upgrade data task should be run immediately before reboot of the HMC with volume 1 of the
recovery CD. If the HMC reboots and does not go to the install menu of volume 1 of the recovery CD, you
should repeat the save upgrade task. The recovery CD is the same as the RG_BASE_040507.1.iso CD. For more
information see:
https://techsupport.services.ibm.com/server/hmc/power4/fixes/mcode/ptf_r3v30mc.html

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

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.

select F8 for New Installation
OR
select F1 for an Upgrade installation.
On the next screen to confirm Select F1.
You will be prompted to insert the second CD.
Remove the CD from the DVD Ram drive and hit enter when the install is completed.

Install the HMC_Update.V3R3.1.zip from the HMC support link:
http://techsupport.services.ibm.com/server/hmc/power4/fixes/mcode/r3v30.html

- Select HMC PTF Update for Service Pack 8
- Use the HMC-->Install Corrective Service option to install.
- Reboot HMC after successful installation.

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

1. HPS Network Topology Reconfiguration Procedures
In Chapter 9, "Service Procedures", of the pSeries High Performance Switch Planning, Installation, and Service Guide, version GA22-7951-01, there is a section titled "Required network cold start". Part of this procedure may require the deletion of the fnmSwConfig file, if you have reduced the number of switches in your network or if you have changed your network topology:
Delete the auxiliary switch topology file fnmSwConfig which can be found in the following directory: /opt/hsc/data
The on-site CE must perform the deletion of that file. If you have not reduced the number of switches in your network or changed your network topology, that file does not need to be deleted.

2. There is a mandatory Bios upgrade required for all Bradley logan HMC PC's. 7315-C03, 7310-C03, 8187-KUH
DownLoad: http://techsupport.services.ibm.com/server/hmc/power4/fixes/mcode/r3v30.html

3. 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 from HMC3.2.6 to HMC3.3.0.

4. wu-ftp will be removed from the HMC distribution. One will be able to ftp out of the HMC but not into the HMC. scp is available if you enable ssh.

5. HMC WebSM PC CLIENT needs to be reinstalled on your remote server or PC in order for WebSM to work with this HMC Build. After you have installed the HMC, use http://<hmc-hostname>/remote_client.html to install the PC Client software on your remote server or PC.

6. 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, pesh command is provided. 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.

HMC for p690

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

For “Code load requirements for existing p690 server frames” please see the section in Chapter 6 titled as such and follow "Step 1: p690 HMC code load" in the HPS guide.

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

Then follow the section in the HPS guide titled “Step 6. Install the Hardware Management Console (HMC)” in chapter 6 and follow the steps till “Step 17. Your System is Now Set Up” in chapter 6 to understand how to connect the rs422/rs232 cables.

Then go to Step 1: p690 HMC code load in chapter 6 and follow the instructions in order to upgrade your HMC to Service Pack 6 and configure the 8-port/128-port adapters on your HMC.

Level Check HMC by doing from the HMC GUI top menu bar task "Help" followed by "About Hardware Management Console" and it should show:
Release 3, Version 3.1 HMC Build Level 20040716.1.

HMC for p655

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

For “Code load requirements for existing p655 server frames” please go the section in Chapter 6 titled as such and follow "Step 1: p655 HMC code load" in the HPS guide.

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

Then follow the section in the HPS guide titled “Step 6. Install the Hardware Management Console (HMC)” in chapter 6 and follow the steps till “Step 17. Your System is Now Set Up” in chapter 6 to understand how to connect the rs422/rs232 cables.

Then go to Step 1: p655 HMC code load in chapter 6 and follow the instructions in order to upgrade your HMC to Service Pack 6 and configure the 8-port/128-port adapters on your HMC.

Level Check HMC by doing from the HMC GUI top menu bar task "Help" followed by "About Hardware Management Console" and it should show:
Release 3, Version 3.1 HMC Build Level 20040716.1.

SNM/FNM for HMC (APAR IY58991, PTF U499536):

Follow the install instructions for this HMC PTF by choosing PTF U499536 on the website http://techsupport.services.ibm.com/server/hmc/power4/fixes/mcode/r3v30.html

Note: Installation of SNM_Serv_Pack8.zip is required on all of the HMCs in the system. The SNM contained and installed as part of the new HMC R3 V3.0 base is not the current level. Additionally, there are some base HMC platform changes that are packaged with SNM_Serv_Pack8.zip that need to be delivered to all of the HMCs. In general, SNM updates are only needed on HMCs that are running SNM; however, this update needs to be installed on all HMCs. Each HMC must be rebooted after installation.

To Level Check this update please follow the same set of web instructions at the bottom titled "Verify a successful update".

HPS/SNI LPPs:

Important: The HPS LPP filesets MUST be installed prior to the GFW driver if the current level is at Sevice Pack 6 or lower. Logical partitions should NOT be rebooted until the GFW driver has been successfully installed on each CEC. Rebooting prematurely will generate "phantom" SNI devices. See Problem #3 in the "Known Problems" section of this document for more information.

NOTE: After successful installation of HPS Filesets from levels at Sevice Pack 6 or lower, technical large page support must be enabled. IBM pSeries HPS now requires that you set up your LPARs with 64 bit kernel and enable Technical Large Page (TLP) option.

Here is a formula to calculate the required TLP (technical large page) used by the SNI support code:

In the formula below, number_of_sni refers to the number of sniX logical interfaces present in the partition. The num_windows, send pool size, and receive pool size values for the AIX partition may be obtained by using the following command:

lsattr -El sniX (where X is the device minor number; 0, 1, 2, etc.)
total_num_windows = num_windows + 7

number of TLP required = A + B + C + D.

where;
A = 1 + (number_of_sni * 2)
B = (number_of_sni * total_num_windows)
C = (number_of_sni * total_num_windows * 262144) / 16777216
D = (send pool size + receive pool size) / 16777216

To setup Large Page option using vmo command below for a node or node group:

vmo -r -o v_pinshm=1 -o lgpg_size=16777216 -o lgpg_regions= <number of TLP required>

OR

dsh -vn <node name> "echo y|vmo -r -o v_pinshm=1 -o lgpg_size=16777216 -o lgpg_regions = <number of TLP required>"

Use the echo command, because vmo will ask for verification to run bosboot
> Setting v_pinshm to 1 in nextboot file
> Setting lgpg_size to 16777216 in nextboot file
> Setting lgpg_regions to the required number of TLP in nextboot file
> Warning: some changes will take effect only after a bosboot and a reboot
> Run bosboot now?
> A previous bosdebug command has changed characteristics of this boot image. Use bosdebug -L to

display what
these changes are.
> bosboot: Boot image is 19877 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.

GFW 3H040602.img (p690) or 3J040602.img (p655)

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.

For each CEC on which you want to install the GFW code -

One partition running AIX must have service authority. Linux does not support microcode download at this time.

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 when the LPAR(s) reactivate.

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

AIX APAR IY49146 is required for update_flash to work correctly. Level Check by running this command on the partitions: " instfix -ik IY49146 "

update_flash 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:

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

To install GFW update from AIX using "update_flash" with a locally available image

1. Disable SNM Software from the SNM GUI Panel.
2. For each CEC on which you want to install the GFW code -
   Shutdown all partitions except the one with service authority
   In the AIX partition with Service Authority:
   Copy the GFW firmware update code(3x040602.img file) to /var
   Where x = H for p690 or J for p655.
   Enter the following command:
   /usr/lpp/diagnostics/bin/update_flash -qf /var/3x040602.img
   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.
   Eg. dsh /usr/lpp/diagnostics/bin/update_flash -qf /var/3x040602.img
3. After the LPAR(s) are ‘Running’, Power OFF the CEC(s) from the GUI
4. Enable SNM Software from the SNM GUI Panel
5. Power ON the CEC(s) from the GUI

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

If you’re installing GFW microcode on a p655 via NIM then follow the HPS guide Chapter 6 under the section titled "Code load requirements for existing p655 server frames", Step 3. GFW (system firmware) code load.

After installing SPCN (which only takes a couple of minutes) power code download will NOT occur on a HPS system containing either a HPS Switch or HPS Adapter(s). Power code for an HPS System is loaded from the HMC so the 2 hour wait period for GFW install does not apply on an HPS system.

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 (Regatta IH series), H=p690 (Regatta H series) yy = year, mm = month, and dd = day of the release.

Check the GFW microcode level from a VTERM to the main SP Menu. This should show the correct level 3x040602 on the top line. 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).......3H040602
```

- OR -

```
ROM Level (alterable).......3J040602
```

Determining the level of HPS adapter microcode (Performed ONLY after Technical Large Page (TLP)support has been enabled - see "HPS/SNI" sub-section in the "Installation Instructions" section below for details on technical large page setup)

- The HPS adapter microcode (ucode) is shipped as part of the GFW update image.
- To Level Check the ucode:
  From AIX partition, issue: /usr/sni/aix52/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. The timestamp there should show it being built on 05/10/04

.

- **NOTE:** 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. Level Check the ucode as stated in the previous bullets to make sure you do not have to reinstall GFW.

Power Subsystem Microcode: ptcode-1.67.2665-1.i386.rpm:

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 4: p690 power subsystem microcode load" or "Step 4: p655 power subsystem microcode code load" respectively in the HPS guide.

- Install via the HMC GUI through the Software Maintenance -> Frame panels.
- Power cycle switch:
  o Using the HMC GUI, select “Switch Network Management <Switch Topology View” option.
  o For each switch plane,
  - select "Selected <Power <Off"
  - refresh GUI to verify power status
  - select "Selected <Power <On"
  - refresh GUI to verify power status
  o Repeat procedure on all switch planes.
- Recycle SNM daemon:
  o Using the HMC GUI, select Switch Network Management<Disable SNM Software/Enable SNM Software.
- Level Check - After completing the ptcode installation, go to the HMC GUI and verify successful instalation.
  - Select: Software Maintenance --> Frame --> Install Corrective Service
  - Verify that the "Installed Version" matches the version you just installed.
  (Important: This window does not automatically refresh when installation completes. Manually refresh the window as necessary.

AIX:

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

- Download and install the recommended AIX 5L version 5.2 Service Level update (IY56769 and IY57276)
as listed in "Section 1" of this document.
- Level Check APAR IY56769 by running this command on the logical partition(s): "lslpp -ha bos.mp64"
and verify that the /usr part of the fileset is at or above 5.2.0.32
- Level check APAR IY57276 by running the following commands on the partitions: " instfix -ik IY57276 "

LPPs:

Document Reference: AIX 5L Version 5.2 Installation Guide and Reference (SC23-4389-03)
Download and install the applicable VSD, LAPI, HPS, PPE, LoadL, GPFS, Parallel ESSL, ESSL, CSM and RSCT PTF updates, as listed in "Section 1" of this document

Section 4: Known problems/workarounds/Restrictions

Problem 1: "Phantom" SNI devices may appear after upgrade from pre-Service Pack 7

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 (3H040602.img or 3J040602.img)
- 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:

) Complete the GFW update.

- 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)

2a) Make note of any customization to the SNI devices (e.g. num_windows, driver_debug, etc ...) using the command "lsattr -El sniX" so that they can be
re-applied after the recovery procedure. The procedure will reset all values back to the defaults.

Remove all the SNI devices:

```
for each logical sniX; # e.g. 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 2b.

- rename the configuration method for the device: mv /usr/sni/aix52/cfgsni /usr/sni/aix52/cfgsni.orig
- reboot each LPAR that failed
- run the rmdev loop again
- restore the original configuration methods name: mv /usr/sni/aix52/cfgsni.orig /usr/sni/aix52/cfgsni

2b) Remove ONLY the top half of the snX devices:
The lower half (original half) have the ipaddr and netmask attributes in the odm. You don't want to delete these, nor do you have to.

```
for UPPER HALF of logical snX # e.g. 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 ...)

Restrictions:

1) 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 Service Pack 8

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.

2) Rules for swapping cables for fault isolation

Component: SNM - Switch Network Management

Systems Affected: All HPS Users applying Service Pack 8

Description:

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.

3) 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, it is recommended that you use the commands below, to recycle and reboot all LPAR/AIX images in your cluster:

"shutdown -F" <-- to shutdown LPAR

"shutdown -Fr" or "dsh -av shutdown -Fr" <-- to reboot LPAR or multiple LPARs simultaneously

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.

5) Do not enable SNM software on more than one HMC at the same time

Component: SNM - Switch Network Management

Description:

After mistakenly starting FNM on two networked connected HMC's, the snX interfaces are not configurable and give error 0821-510.

Solution: If you mistakenly enable SNM software on more than on HMC at the same time, you must disable SNM software on all HMCs where it is enabled, then enable it on a single HMC. If CECs were booting while the multiple instances of SNM software were running, those CECs may need to be rebooted in order to bring that CEC's SNIs into the switch network.

Section 5: HPS Service Pack 8 fix list (by component)

LAPI: Abstract (APAR IY58795)

Inconsistent Tot_counters from new lapi lib
LAPI TLP test case seg faulting with 4G message size
Assertion failed while running Pb xpdgefs on Squadrons

Loadleveler: Abstract (APAR IY58798)

ALTERNATE CENTRAL MANAGER TAKES TO LONG TO TAKEOVER
SETPCREDS ERRMSG RETURNED MAY BE MASKED/INVALID
CONSUMABLECPU AVAILABLE/FREE NOT ACCURATE WHEN CM RETURNS
SCHEDLOG.OLD RETAINING PROCESS PID
llsummary -l -x does not show allocated adapters
total_tasks=odd value fails schedule by preemption
llq -l ! more shows strange characters for Status field
Preempted jobs not resumed due to no machine on server path

Parallel Environment: Abstract (APAR IY58797)

MPI_DUP_INFO() IS FAILING WHEN PASSING A HINT TO MPI_FILE_OPEN
BIND OF 2 OBJ WITH MPI.H CAUSES WARNING MSG
call sites in blocks note available
Large MPI job may hang on machine with many CPUs
64bit MPI job may hang on NH2 Colony system with many CPUs
MPI job failed with No recieve posted for ready mode send
WRONG ALGORITHM CHOICE FOR SMALL MESSAGES
Statement of DPCL support needed
PCT HANGS TARGET APPLICATION - SEGFAULT IN PERFMON_FUNC-R.C
SEGV CONVERTING MPI TRACE RECORD TO SLOG2
Some MPI & MPI LAPI apps core dump when run on LoadLevler

VSD: Abstract (APAR IY58792)

Force installs of rsct.vsd failing
updatevsdnode w/ALL parm fails after problem on 1 adapter
VSD Must verify en1 interface is up prior to put VSD in ACT
allow reads on fenced vsd
vsdnode command failing ERROR::IBM.PeerNode
createvsd fails on with /var full
ON EIO RECOVER VSD.CSER1 SCRIPT SERIALIZES VARYOFFVGS OF VOL
Long waiters on GPFS

GPFS 2.1: Abstract (APAR IY58799)

mmfsck fails with Exit status 0:5:18
mmcrcluster[818]: syntax error at line 819
Man pages problems for mmconfig - AIX cluster
GPFS DSI after mmshutdown
mmcrvsd fails when hdisk doesn't have the pvid
mmcrfs fails to provide message when disk isn't registered
shouldn't be able to register vsd twice
registering rsct vsd makes the switch name it
MMSTARTUP: MMREMOTE: INCORRECT OPTION:
double vsd entry erroneously accepted by mmcrvsd
GPFS CORE FILE FILLED /
LONG WAITER: WAIT FOR BUFFER FETCH
dataStructureDump error message number incorrect
mmcrvsd fails when legitimate duplicate disks
FSCK HANGS AFTER FS ERROR
STUCK IN ::CLOSE AFTER OTHER NODE KILLED
SINGLENODEQUORUM HUNG AFTER ADAPTER FAILURE
panic and assert:(offset < ddbP->mappedLen) file cxiIOBuffer.h
WHEN CLLSIF OUTPUT IS NOT CORRECT, MMCOMMON
mmcrlv does not recover when restarted after failure.
SDR adapter data
False ENOSPC if disk config changes
PROBLEMS WITH VERY LARGE DIRECTORIES
MMRPLDISK DOESNT REACT ON ERRORS POSSIBLY RETURNED FROM TSRPLD
lockSDR perf
deadlock after several recovery scenarios
GPFS.SNAP FAILS WHEN LS ALIASED WITH -X AND ENV FILE
lockSDR variables
LX: Assert exp(dataP->recda.getDiskAddr().isOK())
Better handling of corrupted directories

INVALID DATA IN MMFS.CFG
MEMORY LEAK DOING INODE SCANS
ASSERT ON ENXIO
mmfsck hung
!"Assert on Structure Error", file Logger.C, line 481
HUNG FCNTL LOCK ON AIX 5.2
vsd/nsd deletion progress
error msg with incorrect node name displayed
gpfs.snap missing quote

GPFS 2.2: Abstract (APAR IY58800)

GPFS 2.2 supported with HPS Service Pack 7 or higher

Description:
With the release of HPS Service Pack 7, customers will now be able to use GPFS 2.2 in their HPS environment. The recommend level of service required for GPFS 2.2 on an HPS system is IY58800.
mmchconfig does not always set pagepool
mmchfs -W allowed two file systems with same name
clusterId is a negative number and causes error
mmlspv command
DM_GET_FILEATTR WITH DM_AT_DTIME CAUSING GPFS PANIC
signal 11 in tokem mgr code
signal 11 at: diskSector2AllocSubblk(int,short,Str..,unsigned..
ASSERT: HASREFERENCESP !! GP->GN_FILOCKS==0 VNODEOPS.C
NONROOT USER DENIED ISSUING MMLSQUOTA COMMAND
quiet listNsdPhysicalVolumes
Migrate logfile when replicated disk fails instead of unmount
gcc opt flags for X86_64
Deadlock in file allocators
ASSERT IN DISK::~DISK
Better FG handling when descOnly disk dies
Broken ByteSwap64 on Opteron
Opteron: enable optimization
MMLSFS -D LINES TOO LONG
DR mmfsctl
Consistent IO tracing, signal trace

SNM: Abstract (APAR IY58991)

New SNM function included in HPS Service Pack 7 or higher - Handling of link outages during runtime

Description:

With the SNM package released in HPS Service Pack 7, the route generation algorithm between a source-destination pair (four routes generated), will never be changed.

Implications:

- If a route passes through a faulty link, that route will never be marked "non-preferred" in the path table.
- If one to three routes are marked "non-preferred", packets will be sent using only the remaining preferred routes.
- When all four routes are marked "non-preferred", then the SNI will choose one out of any four routes to send a packet.

Details:

A faulty link includes both a fully broken link or one that has reached a hot error threshold. When an internal link to the network fails, the applications will
experience dropped packets for the duration of time it takes to mark the routes passing through it as "non-preferred" (approximately, 2 minutes).
In general, the applications should not experience performance degradation due to dropped packets.

An exception to this assertion, is the case where a SNI to switch link reaches a hot error threshold and there is a solidly broken link in the path of one of the
four routes between that SNI and some other SNI in the network. Since the SNI to switch link is not fully broken, it can send packets along. When a packet is
sent along the route with the broken link, it will be lost. Hence, the pair will experience a 25% packet loss.

Missing endpoint & riser info for latent errors
ELA gets a Null Pointer Exception when trying to cleanup /var.
Routing wants 5 minute timer for Hot errors
NCAR Extended Error Data
Fix fnm.snap for the restricted shell
Cannot use ps in restricted shell to see fnmd threads
CROSS PLAN PATH UPDATES NOT COMPLETE
Tune CSP IPL_Delay
path table valid bits not restroed
TOD Invalid Error' should not be called home
core dump in routing
Fix extraneous error message in lsSNM
comm sending out-of-order MSGs
NAVO: Incorrect routes between a number of source-destinations
Detected EAGAIN signal when dump link table from i_stub_FS
Checking in new --svcnet option for hps_check.pl

HPS/SNI: Abstract (APAR IY58796)
New CSS function included in HPS Service Pack 7 or higher - Performance enhancements

Description:

SNI device support code has been restructured to deliver higher communication performance in terms of higher bandwidth and lower latency. Most of the restructuring is only surfaced via better performance but some do have implications for the administrator. The number of MPI and LAPI windows is now a tunable parameter with a default of 16 windows and a maximum of 64 windows, The support code allocates and reserves Technical Large Pages (TLP) at boot time and the calculation of the number of pages in the technical large page pool is different. The size of a partition's TLP pool is determined by its
number of SNI devices, size of its IP send and receive pools and the number of windows configured along with the TLP needed for other purposes.

stop building En_US loc filesets.
H_CALL perf timers
dump_key value reset
Test case not getting window fatal notification
panic in ml_output/sn_xmit_spkt
Error data logging for window fatals.
Add dynamic pool reduction in IP support
core dump in hps_snap in handling MP_DOWN

GFW: Abstract
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