

Date : 24/08/2007
Date : 29/11/2007
Date : 30/11/2007
Date : 03/12/2007
Date : 04/12/2007
Date : 07/12/2007
Date : 08/12/2007
Date : 09/12/2007

Version : 0.1
Version : 0.2
Version : 0.3
Version : 0.4
Version : 0.5
Version : 0.6
Version : 0.7
Version : 0.8

Dinil.K.Das
dinildas@in.ibm.com

This document contains extracts from other IBM published open documents.

One of the appreciable aspects of AIX Operating system is that you can customize or change almost all subsystems after the installation other than Trusted Computing Base and file system type. That too, the modifications are mostly on the fly or can be put into effect just with a reboot. So in this discussion we are going to concentrate on the best practices of AIX configuration. It is more on the base practices than best practices.

Topics discussed & planned to include later are as given below.

- | | |
|--|-----------------------|
| 1) <u>TIMEZONE, TIME & DATE.</u> | <u>Done</u> |
| 2) <u>Mirroring rootvg disks.</u> | <u>Done</u> |
| 3) <u>Recommended additional packages.</u> | <u>Done</u> |
| 4) <u>File system structure.</u> | <u>Done</u> |
| 5) <u>AIX Update Strategies: MLs, TLs & SPs.</u> | <u>Done</u> |
| 6) <u>System Firmware Upgrade.</u> | <u>Done</u> |
| 7) <u>IO Firmware Upgrade.</u> | <u>Done</u> |
| 8) <u>AIX Tuning – System & User Environment.</u> | <u>Done</u> |
| 9) <u>Driver Installation / Upgrades.</u> | <u>Done</u> |
| 10) <u>A brief hardening on AIX.</u> | <u>Planned in 1.x</u> |
| 11) <u>pSeries data collection</u> | <u>Planned in 1.x</u> |
| 12) <u>AIX data collection : (snap & perfpmr).</u> | <u>Done</u> |
| 13) <u>Configuring ESA</u> | <u>Planned in 1.x</u> |

Note : It is recommended to discard the factory installation of AIX and reinstall the system freshly with only the required bos packages and suggested additional packages.

TIMEZONE, TIME & DATE.

First and the foremost is to set the time and date in system. Timezone change to the local geographical zone is a suggested practice and must be done for some setups where precise synchronization is expected.

AIX calculates Timezone by the formula

Time zone offset = Co-ordinated Universal Time (CUT) – Local Time

CUT = UTC = GMT all mention the same reference time co-ordinate.

This make India having a Timezone offset of “-5:30” hours.

This is the reverse calculation done compared to some of the other operating systems.

To Change Time Zone, use : *smitty system*
 Select : *Change Time Zone Using User Inputted Values*
 Key in your : *Standard Time ID(only alphabets) :*

In India we follow *IST (Indian Standard Time)*. No need of Day Light Savings settings for India. You can customize it for you country or geographical area.

AIX defines an environment variable *TZ* in the */etc/environment* file.

An AIX user can check the *TZ* variable value by : *echo \$TZ* command.

User can also modify this variable to reflect their local time, in case if it is different from the system's local time. User's profile of the shell will be useful in case the change needs to be permanent.

We all know how to see time & date and how to set it. The date command or *smitty date* , to change time, simple enough. But in a data center or in an organization it is equally important to have the systems running the same time, otherwise they all should be time synchronized. For that, *time synchronization* need to be configured in AIX. Moreover for clustering configuration, time synchronization among the cluster nodes are very important. There are 2 methods. First method is by running the *timed daemon* and the second one is to run *ntp daemon*. Second is more preferred since ntp can work across ip subnets, more precise, hierarchically manageable, widely used, heterogeneous design and after all, the standard in Internet. So our further discussion is limited on to the ntp configuration.

Find out your customer is running any time server in the network, of course which is ntp compatible. If yes, configure AIX systems as ntp client for the time server, else configure one of the AIX servers as the ntp server.

Configure [AIX as ntp client](#) :

Edit : `“/etc/ntp.conf”`
 Add a line : `“server <ntpserver>”`
 where ntpserver = *ip-address / hostname* of the ntp server.

Multiple ntp server entries can be added into the file /etc/ntp.conf for redundancy.

Start ntp daemon : `startsrc -s xntpd`
 For running xntpd persistent across reboot use : `smitty xntpd`
 Stop ntp daemon : `stopsrc -s xntpd`
 Manually Synchronize time : `ntpdate <ntpserver>`
 Check Synchronization status : `lssrc -ls xntpd`

Service xntpd may take a long time to synchronize with the server. May be around 15min to 20min sometimes. If the time difference between the server and client is more, then xntpd may not be able to automatically synchronize time. In that case, first stop xntpd with `stopsrc -s xntpd` and then use the command `“ntpdate <ntpserver>”` to manually synchronize the time with the time server. When you execute this command look for the offset value near (practically this will never be exact) to numeric `“0”`. Run again if required. Then start the daemon with `“startsrc -s xntpd”`.

Check the time synchronization with `“lssrc -ls xntpd”`. Ensure that the parameter `“Sys stratum”` is not `“16”`. Value `“16”` means the time is not synchronized. So expect a value lower than 16, typically between 3 to 6. Client's stratum is generally one value lower (numerical value higher) than the server. Output also shows the ntp server as the value for the parameter `“Peer”`.

Configure [AIX as ntp server](#) :

Edit the file : `/etc/ntp.conf.`
 Comment out the option : `broadcastclient`
 Add the line : `“server 127.127.1.x”`
 where value of `“x”` makes the server to run at stratum level `x`.
 Typically value of x will be 3, 4 or 5.

Then start the ntp service & check the status as given below.

Start ntp daemon : `startsrc -s xntpd`
 For running xntpd persistent across reboot use : `smitty xntpd`
 Stop ntp daemon : `stopsrc -s xntpd`
 Check Synchronization status : `lssrc -ls xntpd`

Then other servers can synchronize with this ntp server.

Check the time synchronization daemon status with `“lssrc -ls xntpd”`. Ensure that the parameter `“Sys stratum”` is not `“16”`. Value `“16”` means the time is not synchronized. So expect the value `“x”` with is the stratum level defined in the `/etc/ntp.conf` file.

```
root@bgl-wks-con-09:~# lssrc -ls xntpd
Program name:      /usr/sbin/xntpd
Version:           3
Leap indicator:    00 (No leap second today.)
Sys peer:          bgl-svr-dns-03.bsnl.net.in
Sys stratum:       4
Sys precision:     -18
Debug/Tracing:    DISABLED
Root distance:     0.284119
Root dispersion:   0.259644
Reference ID:      10.16.17.13
Reference time:    caf8295e.9a3c7000  Wed, Nov 28 2007 23:22:30.602
Broadcast delay:   0.003906 (sec)
Auth delay:        0.000122 (sec)
System flags:      pll monitor filegen
System uptime:     526477 (sec)
Clock stability:   24.798904 (sec)
Clock frequency:   0.000000 (sec)
Peer: bgl-svr-dns-03.bsnl.net.in
  flags: (configured)(sys peer)
  stratum: 3, version: 3
  our mode: client, his mode: server
Peer: bgl-svr-dns-04.bsnl.net.in
  flags: (configured)(sys peer)
  stratum: 3, version: 3
  our mode: client, his mode: server
Subsystem          Group          PID          Status
xntpd              tcpip          741548       active
#
```

Mirroring rootvg disks.

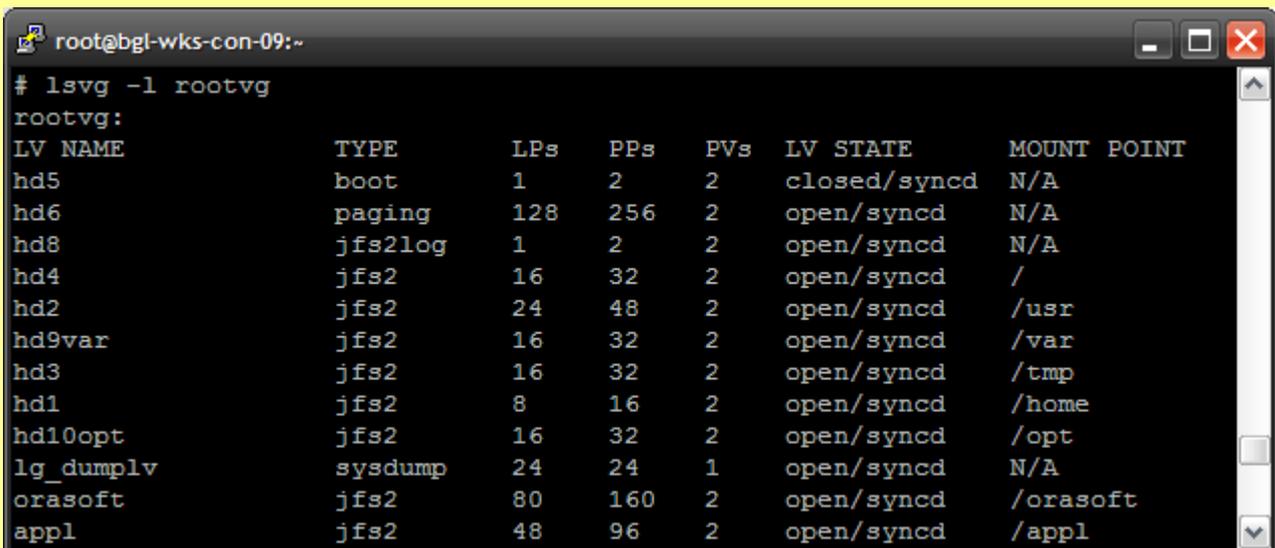
There is no need to emphasize the need for mirroring rootvg in AIX. After all who wanted to restore the mksysb backup for a simple reason for rootvg disk failure. Mirroring rootvg in AIX is a very simple & straight task. You can attempt mirroring before increasing the file system size as it takes comparatively less amount of time.

```
extendvg rootvg <hdisky>
mirrorvg -m rootvg <hdisky>
bosboot -a
bootlist -m normal <hdiskx> <hdisky>
```

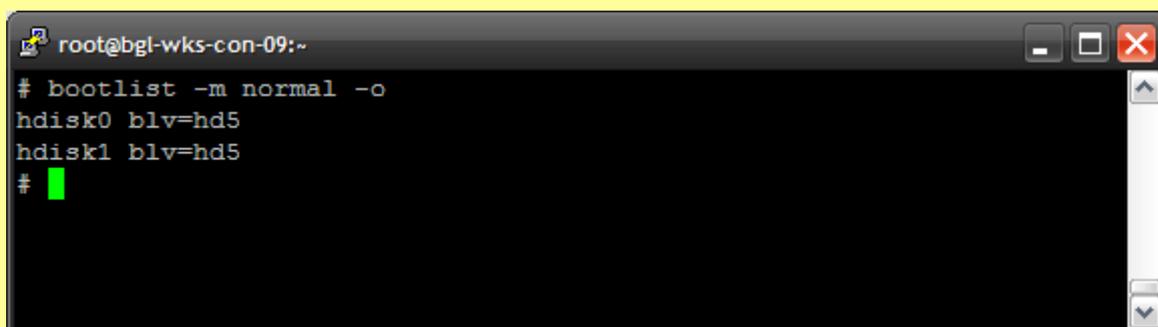
where hdiskx is the hdisk where rootvg is on and hdisky is the mirrored disk of hdiskx.

Please do remember to have the paging space mirrored. Though mirroring paging space is a performance problem, a system without mirrored paging space may crash when the disk having paging space fails.

Ensure that mirroring is done as follows



```
root@bgl-wks-con-09:~
# lsvg -l rootvg
rootvg:
LV NAME      TYPE      LPs      PPs      PVs      LV STATE    MOUNT POINT
hd5          boot      1        2        2        closed/syncd  N/A
hd6          paging    128      256      2        open/syncd    N/A
hd8          jfs2log   1        2        2        open/syncd    N/A
hd4          jfs2      16       32       2        open/syncd    /
hd2          jfs2      24       48       2        open/syncd    /usr
hd9var       jfs2      16       32       2        open/syncd    /var
hd3          jfs2      16       32       2        open/syncd    /tmp
hd1          jfs2      8        16       2        open/syncd    /home
hd10opt      jfs2      16       32       2        open/syncd    /opt
lg_dump1v   sysdump   24       24       1        open/syncd    N/A
orasoft     jfs2      80       160      2        open/syncd    /orasoft
appl        jfs2      48       96       2        open/syncd    /appl
```



```
root@bgl-wks-con-09:~
# bootlist -m normal -o
hdisk0 blv=hd5
hdisk1 blv=hd5
#
```

```
root@bgl-wks-con-09:~# lslv -l hd5
hd5:N/A
PV          COPIES          IN BAND          DISTRIBUTION
hdisk0      001:000:000      100%             001:000:000:000:000
hdisk1      001:000:000      100%             001:000:000:000:000
# lslv -l hd6
hd6:N/A
PV          COPIES          IN BAND          DISTRIBUTION
hdisk0      128:000:000      100%             007:093:028:000:000
hdisk1      128:000:000      100%             000:109:019:000:000
# lslv -l hd8
hd8:N/A
PV          COPIES          IN BAND          DISTRIBUTION
hdisk0      001:000:000      100%             000:000:001:000:000
hdisk1      001:000:000      100%             000:000:001:000:000
# lslv -l hd4
hd4:/
PV          COPIES          IN BAND          DISTRIBUTION
hdisk0      016:000:000      50%              000:000:008:008:000
hdisk1      016:000:000      50%              000:000:008:008:000
# lslv -l hd2
hd2:/usr
PV          COPIES          IN BAND          DISTRIBUTION
hdisk0      024:000:000      66%              000:000:016:008:000
hdisk1      024:000:000      66%              000:000:016:008:000
#
```

Installing additional packages.

Though the AIX default installation is comprehensive enough to address most of the package requirements, due to various reasons additional packages need to be installed into the system. Following are some of them which are preferable.

From AIX LPP CDs

<i>X11.Dt</i> <i>X11.adt</i> <i>X11.apps</i> <i>X11.base</i> <i>X11.compat</i> <i>X11.man.en_US</i> <i>X11.motif</i> <i>X11.help.en_US</i> <i>X11.vfb</i> <i>X11.vsm</i>	<i>Java</i> <i>bos.adt</i> <i>bos.cifs_fs</i> <i>bos.data</i> <i>bos.perf</i>
---	---

From AIX ToolBox

bash
bash-doc
lsof
openssl
openssl-devel
openssl-doc
rsync
sudo
unzip
zip

From AIX Expansion Pack

Java 64bit
openssh.base
openssh.license
openssh.man.en_US

File System Structure.

AIX file systems.

We all know that AIX installation automatically selects the file system size as required. File system type jfs2 is a good suggestion at the time of installation (this is the default now). After installation is over, file system size can be increased on the fly. Below given are some recommended file system sizes considering the AIX operating system under normal operation in most of the cases. May be depending on the application requirements the file system sizes can be increased further. After all, it is an online activity and on AIX 5.3, even the shrinking of jfs2 file systems can also be done online.

/	1GB
/usr	4GB
/var	4GB
/tmp	4GB
/opt	1GB
/home	1GB

Paging space.

The common practice of keeping paging space as twice the RAM is a good suggestion up to around 8GB of real system memory. Beyond this, size of paging space can be around 16GB unless otherwise application requirements drive it further up. Paging space is not an alternate to main memory. Always remember to keep the paging space under mirroring or on a disk backed up by any hardware redundant RAID technology because the system may crash if the page space becomes inaccessible. It is definitely a good practice to spread the paging space on multiple spindles for performance. Considering flexibility also it is also a good practice to create multiple paging space. For eg: if a 16GB paging space is the requirement consider creating 2 paging space of 8GB each, and position them on different set of spindle pairs. Please also remember not to delete the system default page space usually hd6 in the rootvg. Removing default paging spaces incorrectly can prevent the system from restarting.

Keep page space twice as RAM if RAM <= 8GB.

Keep page space at 16GB if applications are not demanding more.

Always have paging space backed up by hardware redundancy technology.

Spread the paging space on multiple disk spindles for performance.

Divide the paging space into multiple paging spaces for greater flexibility.

Do not delete the system default paging space hd6. (Follow special instructions if it is must)

Do not reduce the size of hd6 below 32MB.

All paging space should be of roughly equal size.

Do not extend a single paging space to multiple physical volumes.

Select spindles with least activity

If possible use multiple disk controllers.

System Dump space

By default the dump devices defined in the SWservAt ODM object class are used. The default primary dump device is `/dev/hd6`. The default secondary dump device is `/dev/sysdumpnull`.

The default configuration in AIX 5.2 is as follows :

```
# sysdumpdev -l
primary          /dev/lg_dumplv
secondary        /dev/sysdumpnull
copy directory   /var/adm/ras
forced copy flag TRUE
always allow dump FALSE
dump compression OFF
```

The default configuration in AIX 5.3 is :

```
# sysdumpdev -l
primary          /dev/hd6
secondary        /dev/sysdumpnull
copy directory   /var/adm/ras
forced copy flag TRUE
always allow dump FALSE
dump compression ON
```

AIX does not support dumping to a mirrored logical volume. A mirrored paging space may be used as a dump device. So there is no worry of creating a second dump lv on the mirrored disk. But paging space holds the dump only for one time. That is as soon as a dump is detected, system will copy that dump to the copy directory. But if that fails first time, dump is lost if it is on paging space. If the dump is on a separate lv, the lv can hold the dump till the next dump occurs making it flexible for administrators to copy the dump later.

Compressing dumps reduces the size needed for dump devices, but may cause the dump process to take longer.

“sysdumdev -c” to turn OFF dump compression.

“sysdumpdev -C” to turn ON dump compression. Ratio is around 1:5 times.

“sysdumpdev -e” command to find out the size requirements of dump dev.

The primary dump devices must always be in the root volume group for permanent dump devices. A suggestion is to keep the following configuration.

```
# sysdumpdev -l
primary          /dev/lg_dumplv
secondary        /dev/sysdumpnull
copy directory   /var/adm/ras
forced copy flag TRUE
always allow dump FALSE
dump compression ON
```

Using lppchk : Ensuring the Operating System integrity.

The lppchk command verifies that files for an installable software product (fileset) match the Software Vital Product Data (SWVPD) database information for file sizes, checksum values, or symbolic links.

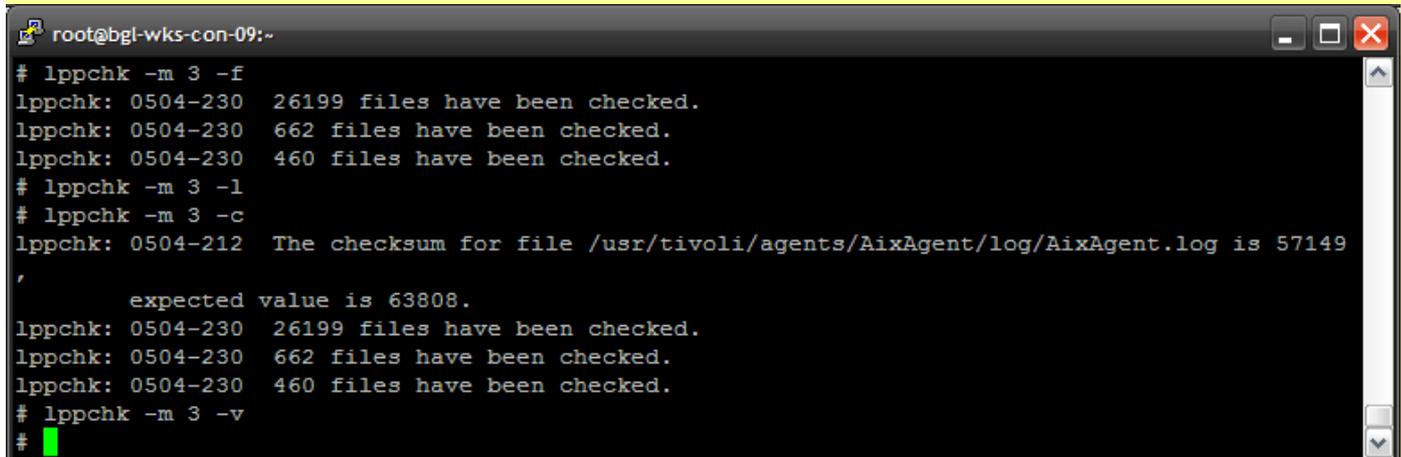
We prefer to check all the 4 types of checks for ensuring Operating System integrity.

lppchk -m 3 -f : Checks the FileList items are present & the file size matches the SWVPD database.

lppchk -m 3 -l : Verifies symbolic links for files as specified in the SWVPD database.

lppchk -m 3 -c : Performs a checksum operation on the FileList items & verifies that the checksum and the file size are consistent with the SWVPD database.

lppchk -m 3 -v : Verifies that the / (root), /usr and /usr/share parts of the system are valid with each other.



```

root@bgl-wks-con-09:~# lppchk -m 3 -f
lppchk: 0504-230 26199 files have been checked.
lppchk: 0504-230 662 files have been checked.
lppchk: 0504-230 460 files have been checked.
# lppchk -m 3 -l
# lppchk -m 3 -c
lppchk: 0504-212 The checksum for file /usr/tivoli/agents/AixAgent/log/AixAgent.log is 57149
,
expected value is 63808.
lppchk: 0504-230 26199 files have been checked.
lppchk: 0504-230 662 files have been checked.
lppchk: 0504-230 460 files have been checked.
# lppchk -m 3 -v
#
  
```

This is a very good tool to check the Operating System consistency. So many times this helps us to find out mistakes in the system which manifests as command failures or system mis-behaviors.

lppchk -u : Can be used for updating the new checksum or size information from the system when the system information does not match the SWVPD database. This flag sets symbolic links that are found to be missing.

AIX Update Strategies : MLs, TLs & SPs.

In the field there are **3 strategies** of updating AIX Operating System prevailing in the field.

- 1) Maintenance Level (ML) updates.
- 2) Technology Level (TL), Service Pack (SP) ending with Concluding Service Pack (CSP) update.
- 3) Technology Level updates (TL) & Service Pack (SP) updates for 2 years (The new support strategy)
This strategy is supporting mainly AIX 5.3 with TL6 onwards.

For years AIX was following the strategy of providing regular updates to the Operating System as Maintenance Levels (ML). This was the practice till [AIX 5.2 ML06](#) and [AIX 5.3 ML03](#). IBM releases Maintenance Levels approximately every half yearly. Maintenance Levels are all inclusive bundles with enhancements and patches.

The AIX 5L OS Release and Service Strategy for 2006

Starting in 2006, a Maintenance Level will be referred to as a Technology Level and will only be released twice per year. Having a year between major releases of new software features also allows for more extensive testing of these releases.

For clients who wish to install only one major level per year, the concepts of a Service Pack and Concluding Service Pack will be introduced. A Concluding Service Pack will allow support on a Technology Level for up to fourteen months from the date the Technology Level was released.

This strategy is applicable for :

AIX 5.2 TL07, AIX 5.2 TL08 & AIX 5.2 TL09
AIX 5.3 TL04 & AIX 5.3 TL05

Technology Level (TL)

A Technology Level is the new term for the twice yearly releases which contain new hardware and software features, and service updates. The first TL will be restricted to hardware features and enablement, as well as software service. The second TL will include new hardware features and enablement, software service, and new software features, making it the larger of the two yearly releases. In the past, the term Maintenance Level (ML) even though the release contained code for new features and functions. Technology Level is a more appropriate term for this code.

The Technology Level (e.g., 5300-04) will be displayed using the “*oslevel -r*” AIX 5L command, just as this command can currently be used to display the Maintenance Level.

Installing a Technology Level should be viewed as an “all or nothing” operation, meaning that requisites will be added so that the whole Technology Level is installed, and not allow a TL to be partially installed. This change is to avoid getting into unsupported configurations.

Rejecting a Technology Level once it has been applied to an AIX 5L system is not recommended. Since Technology Levels are usually large, it is faster and less risky to fall back to the previous level using other methods. To fall back with a reboot, use `alt_disk_install` or the new `multibos` function shipped with 5300-03. To fall back with a restore, create a backup (`mksysb`) to NIM or bootable media (CD, DVD, or tape) before applying the Technology Level.

Service Pack (SP)

The Service Pack concept will allow service-only updates (as known as PTF's) that are released between Technology Levels to be grouped together for easier identification. These fixes will be for highly pervasive, critical, or security related issues. Service Packs will be provided for the N and N-1 releases (e.g., V5.3 and V5.2) on the latest Technology Level for each release (e.g. 5300-04 and 5200-08), and they will be released approximately every 4-6 weeks after the release of a Technology Level.

Applying and rejecting an individual service update (PTF) is still a supported and recommended method of removing an update if there is a problem or a regression after it is installed. Since Service Packs can also be rejected, it is recommended that before applying a Service Pack or PTF update, that all other updates on the system are committed.

Applying the latest level of available updates will move the system to the latest Service Pack. To see which Service Pack is currently installed, run the “*oslevel -s*” command. Sample output for a V5.3 system, with Technology Level 4, and Service Pack 2 installed would be:

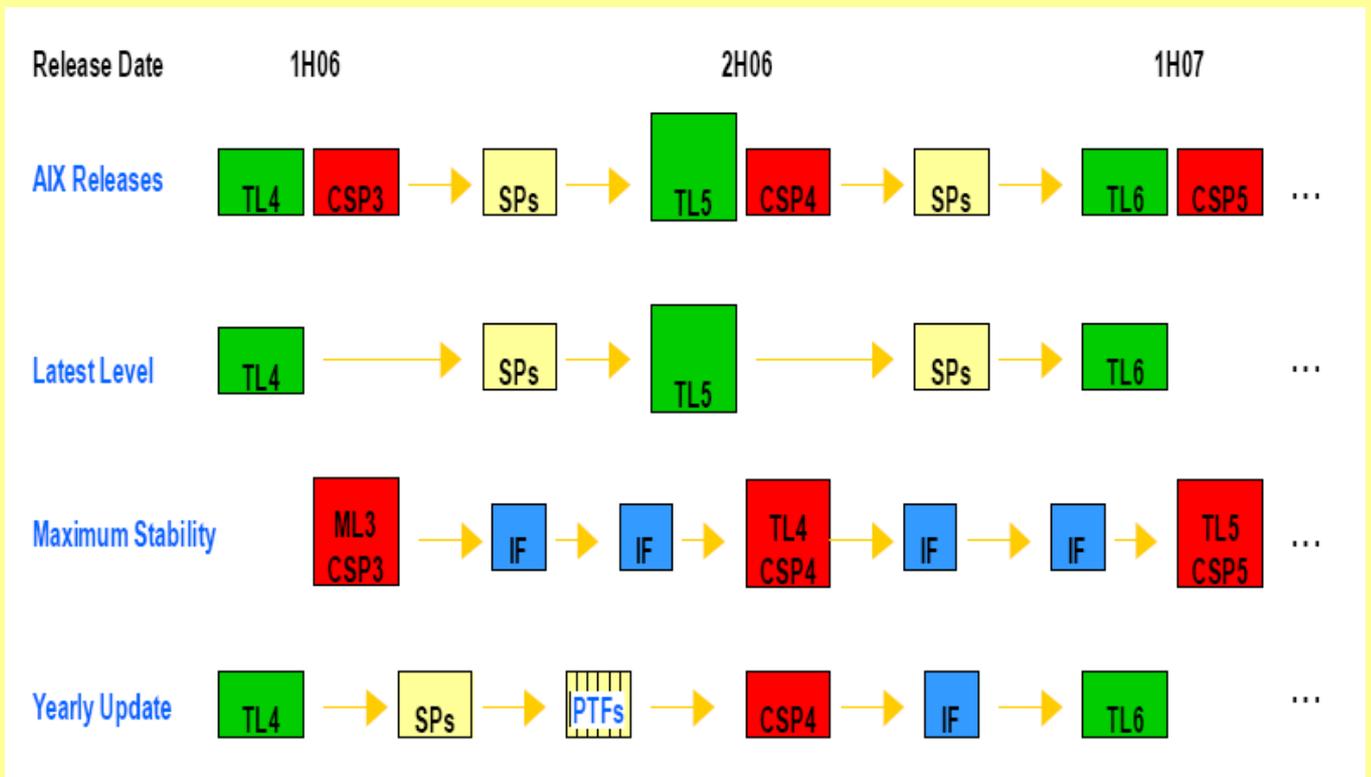
```
oslevel -s
5300-04-02
```

Service Packs are cumulative, so if Service Pack 3 is applied, all of the previous critical fixes from Service Packs 1 and 2 will also be applied. “Critical Fix Packs” will no longer be shipped, since all of the updates that are contained in a Service Pack have been deemed critical or pervasive.

Concluding Service Pack (CSP)

Concluding Service Pack is the term that will identify the last Service Pack on a Technology Level. The CSP will contain fixes for highly pervasive, critical, or security related issues, just like an SP, but it may also contain fixes from the newly released Technology Level that fall into these categories. Therefore, a CSP will contain a very small subset of service that was just released as a part of a new Technology Level.

The Concluding Service Pack will be available shortly after a new Technology Level is released. For example (dates may change), if Technology Level 5300-04 is released in February of 2006, the CSP for the previous release, 5300-03, will be available approximately 4-8 weeks later. It will have a specific level designation of “CSP”. For example, the output of running the command “*oslevel -s*” would return “5300-03-CSP”. Concluding Service Packs will allow for extended service on a Technology Level through the utilization of Interim Fixes.



The AIX 5L OS Release and Service Strategy from 2007

IBM is enhancing the AIX 5L™ operating system (OS) Release and Service Strategy in 2007 as part of the ongoing effort to improve the manageability and stability of the AIX 5L operating system for our clients. The enhanced strategy will provide clients with:

- ∇ Longer support for each AIX 5L Technology Level update (formerly known as MLs)
- ∇ Improved serviceability for the AIX 5L OS throughout the life of each Technology Level
- ∇ Support for some new hardware on previous Technology Levels

The new service strategy will only be used in **AIX 5L Version 5.3 OS**, starting with **TL6**. The CSP for TL5 will be released soon after TL6 is released and interim fix support will still be offered on TL5 for six months. AIX 5L Version 5.2 OS not use the new Service Strategy. When the final TL level is released, new SP's will be released until it goes end-of-life.

Contents of a Service Pack : Service packs contain fixes for:

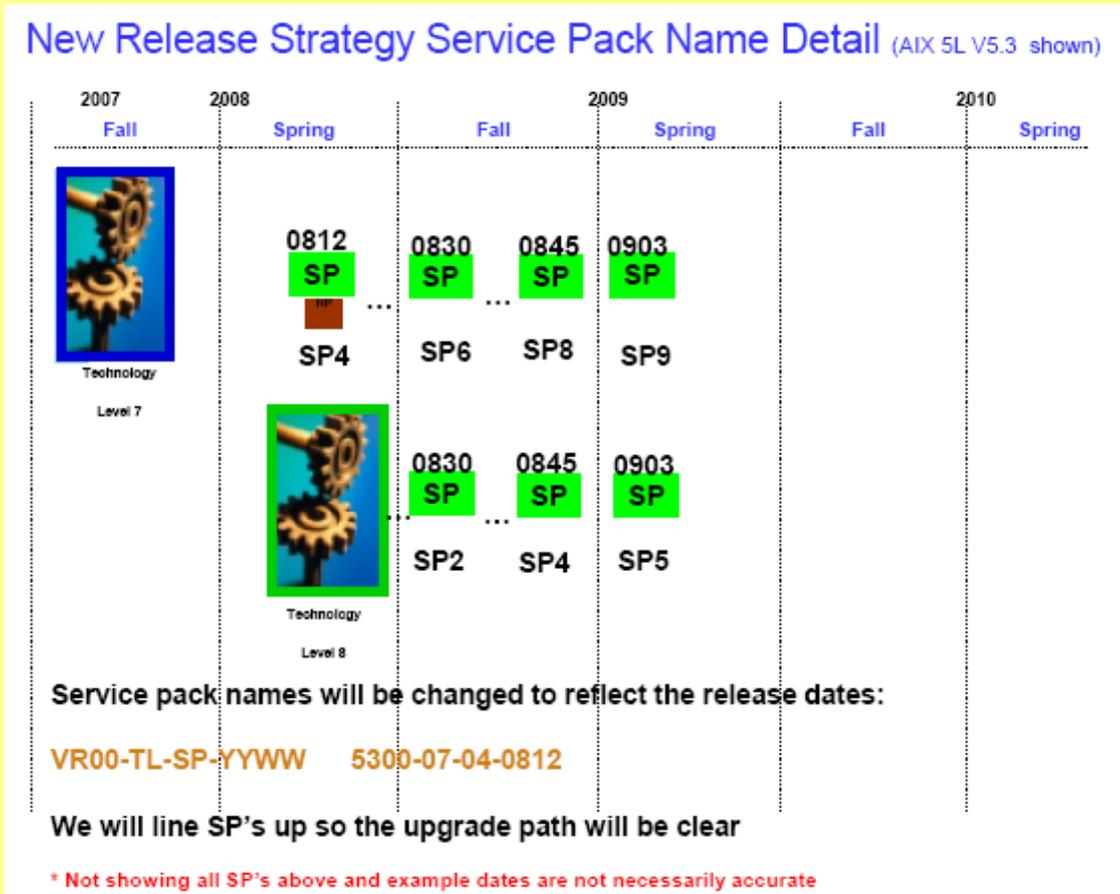
- ∇ Customer reported problems (APARs) that cannot wait until the next TL
- ∇ Critical problems found by development or test teams
- ∇ Very, very limited number of changes to support new hardware.

Examples: New device drivers, new ODM entry to allow for configuration of a new class or type of device, small changes in kernel to recognize a new processor speed, etc.

The only changes that are allowed in an SP are limited to corrections that do not change behavior, cause regression or add new functionality. New function, both for hardware exploitation and software features, is only shipped in Technology Levels or new releases.

Changing the names of the service packs (adding release dates)

Starting with TL6, Service Packs will be labeled with their release date, using the YYWW format, where MM is the 2-digit month and WW is the 2 digit week the SP is available. For example, if 5.3 TL6 SP3 was released in the first week in September 2008, it would be called 5300-06-03-0836 (oslevel -s will report this new level name). The nice thing about this is not only can you tell when the SP was released, but you will also know which SP to move to on a new TL. When moving up to a new TL, you must move to a SP that is the same or later than your current SP. The SP number itself will not be the same, because the Service Packs will be numbered consecutively as they are released, but the dates will tell you where you need to be on the new TL.



The rule of thumb should be that when moving to a new TL, move to the latest SP, that way you are guaranteed it will install. The installp command was changed, as were the updates themselves, to not allow a system to apply any updates that are 'younger' than what is currently installed. This would cause regression.

Each fileset update in a SP has the date marked (just like the SP) and the installp command has been changed to not allow an old SP/fileset on top of a newer SP/fileset, even if that SP is on an older TL. So, using the example above, if you are currently at **TL7 SP8 (5300-07-08-0845)**, then you would not be allowed to install **TL8 SP2 (5300-08-02-0830)**, because 0845 was released after 0830, and therefore will have fixes and hardware enablement that 0830 does not have.

The oslevel command will now print out the new SP format, but the other options will not change.

```
# oslevel
5300
# oslevel -r
5300-07
# oslevel -s
5300-07-02-0811
```

The `installp` command will also stop an update entirely if it sees any updates that are older (released before) in the list it is trying to apply. This is to make sure that part of a SP is not installed, if that was not your intention. If you see a message from `installp` about not being able to install a fileset update because of regression, then go to a newer SP (the latest) and try the install again.

Service Pack schedules

Because IBM will be supporting more TL's in the field and because IBM want to line the SP's up to come out at the same time (with the same YYWW), then SP's release schedules will be increased to approximately every 8-12 weeks. Occasionally, if a critical problem is found (security, for example) and we cannot wait for the next SP to release the fix, a new service pack will be released that will only have one or maybe two changes different from the previous SP. But, it will get a new SP number, as well as a new release date. Depending on where the problem occurred will depend on which TL's get a new SP.

mksysb Backups

As always, you should create backups (`mksysbs`) of your system before and after any update or upgrade. If backing up to media (CD, DVD or tape), a boot image will be created for you on that media.

NIM mksysb

If you are using NIM, be sure to use the `mksysb` image as a source for creating your SPOT/boot image, since that will guarantee that your boot image will match your `mksysb`. The option to create a SPOT/boot image from the `mksysb` image has been available since version 5.3.

Moving to a New TL

You should move to a new TL:

- ∇ If your existing TL is out or is about to go out of service
- ∇ You want to use new function and/or features in a new TL. Hardware exploitation, such as large page space or new software function, such as `multibos`, will only be released in a TL.
- ∇ You are going to test a new level and want to get the longest support possible. In this case, you should move to the latest TL.

Moving to a New SP

If you are currently running on a supported TL, then any SP will be supported, but the risk in staying at an older SP is being asked to move up to a newer SP because it contains an update that you need. Updating twice a year to a new SP or TL is recommended to stay current.

Under any strategies, it is very easy to upgrade the AIX patches. All of them MLs, TLs or SPs can be updated using the smitty option : *smitty update_all*.

```

Update Installed Software to Latest Level (Update All)

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                     [Entry Fields]
* INPUT device / directory for software      .
* SOFTWARE to update                          update_all
PREVIEW only? (update operation will NOT occur)  no      +
COMMIT software updates?                      yes      +
SAVE replaced files?                          no       +
AUTOMATICALLY install requisite software?     yes      +
EXTEND file systems if space needed?          yes      +
VERIFY install and check file sizes?         no       +
DETAILED output?                             no       +
Process multiple volumes?                    yes      +
ACCEPT new license agreements?               no       +
Preview new LICENSE agreements?              no       +

F1=Help      F2=Refresh      F3=Cancel      F4=List
F5=Reset     F6=Command     F7=Edit       F8=Image
F9=Shell     F10=Exit       Enter=Do

```

Be sure to try a preview and accept license as well.

[Package installation](#)

Follow these steps to install the update package:

1. Always run the inutoc command to ensure the installation subsystem will recognize the new fix packages you download. This command creates a new .toc file for the fix package. Run the inutoc command in the same directory where you downloaded the package filesets. For example :

```
inutoc /usr/sys/inst.images
```

2. To rename the downloaded files to their fileset name, run the bffcreate command. After renaming the files, run the inutoc command again. For example, if you downloaded the filesets to /usr/sys/inst.images, run the following command to rename them:

```
bffcreate -c -d /usr/sys/inst.images
```

3. For selected updates

To install selected updates from this package, use the following command:

```
smit update_by_fix
```

4. For all updates

To install all updates from this package that apply to the installed filesets on your system, use the following command: *smit update_all*

It is highly recommended that you apply all updates from the package.

5. Reboot the system : A reboot is required for this update to take effect. *shutdown -Fr now*

Fix packs

: for AIX 5.2 operating system

Technology Levels and Service Packs			
Name	Type	Prereqs	Date
5200-10-03-0744	Service Pack	TL 5200-10	October 2007
5200-10-02-0730	Service Pack	TL 5200-10	July 2007
5200-10-01-0722	Service Pack	TL 5200-10	June 2007
5200-10	Technology Level		June 2007
5200-09-CSP	Concluding Service Pack	TL 5200-09	June 2007
5200-09-06	Service Pack	TL 5200-09	March 2007
5200-09-05	Service Pack	TL 5200-09	February 2007
5200-09-04	Service Pack	TL 5200-09	January 2007
5200-09-03	Service Pack	TL 5200-09	November 2006
5200-09-02	Service Pack	TL 5200-09	October 2006
5200-09-01	Service Pack	TL 5200-09	September 2006
5200-09	Technology Level		August 2006
5200-08-CSP	Concluding Service Pack	TL 5200-08	August 2006
5200-08-02	Service Pack	TL 5200-08	April 2006
5200-08-01	Service Pack	TL 5200-08	February 2006
5200-08	Technology Level		February 2006
5200-07-CSP	Concluding Service Pack	5200-07	March 2006

Maintenance Level packages (legacy)			
Name	Type	Prereqs	Date
5200-07	Maintenance Level		September 2005
5200-06	Maintenance Level		May 2005
5200-05	Maintenance Level		January 2005
5200-04	Maintenance Level		December 2004
5200-03	Maintenance Level		May 2004
5200-02	Maintenance Level		October 2003
5200-01	Maintenance Level		May 2003

: for AIX 5.3 operating system

Fix packs			
Name	Type	Prereqs	Date
5300-07-01-0748	Service Pack	5300-07-00-0747	November 2007
5300-07-00-0747	Technology Level		November 2007
5300-06-04-0748	Service Pack	5300-06	November 2007
5300-06-03-0732	Service Pack	5300-06	August 2007
5300-06-02-0727	Service Pack	5300-06	July 2007
5300-06-01-0722	Service Pack	5300-06	June 2007
5300-06	Technology Level		June 2007
5300-05-CSP	Concluding Service Pack	5300-05	May 2007
5300-05-06	Service Pack	5300-05	March 2007
5300-05-05	Service Pack	5300-05	February 2007
5300-05-04	Service Pack	5300-05	December 2006
5300-05-03	Service Pack	5300-05	November 2006
5300-05-02	Service Pack	5300-05	October 2006
5300-05-01	Service Pack	5300-05	September 2006
5300-05	Technology Level		August 2006
5300-04-CSP	Concluding Service Pack	5300-04	August 2006
5300-04-03	Service Pack	5300-04	May 2006
5300-04-02	Service Pack	5300-04	April 2006
5300-04-01	Service Pack	5300-04	March 2006
5300-04	Technology Level		February 2006
5300-03	Maintenance Level		September 2005
5300-02	Maintenance Level		May 2005
5300-01	Maintenance Level		January 2005

System Firmware Upgrade.

As a suggested practice we must upgrade the System Firmware to the latest one. Too latest is not my general preference. I suggest for a firmware which is with a minimum GA release of 1 month. Please do have a copy of the new firmware release notes. This is a very useful information.

There are links to all pSeries sites from the parent site : <http://www.ibm.com/support/>

System Firmwares at the time of this writing are :

for **p5** machines is **01SF240_332_201** including Bulk Power **02BP240_219_168** firmware.
for **p6** machines is **EM310_057_048**

For Engineers we need to maintain only one copy of the firmware CD which works across all the p5 machines. This is made from the iso image and used with HMC for Upgrade.

A typical upgrade on a single frame p595 or p590 will take around 1.5 Hours.
and on a multi frame machine may take around 2.5-3 hours.

On a p575 frame (one p575+frame) it may take around 1 hour.

A typical upgrade on a fully configured p570 may take around 40-45 minutes.
For other low end and entry level machines also plan for around 40-45 minutes of upgrade time.

If current version of the running microcode is above 01SF240_219, then mostly the updates can be done concurrently. But if the current version is below 01SF240_284, a reboot of the Managed System is required for some deferred updates to get activated.

HMC allows to upgrade multiple machines together with the same firmware. But these Managed Systems should belong to different frames. This is because the HMC puts a lock on the frame during firmware is getting upgraded.

POWER code matrix

Latest release levels for IBM System i and System p

Supported code combinations for HMC and server firmware

Supported HMC and Server code combinations : Excluding 590 and 595

The following table lists currently supported firmware (FW) Release Levels for POWER5 systems, as well as the compatibility of HMC FW levels with system FW levels.

HMC levels	POWER 5 system firmware levels (iSeries and pSeries)			
	240 Release	235 Release	230 Release	225 Release
HMC V7R3 Minimum HMC level to support POWER6. Also supports POWER5 systems at Release Level SF240_299 and above	Supported combination for Service Pack SF240_299 and higher	Not a recommended or supported combination	Not a recommended or supported combination	Not a recommended or supported combination
HMC V6R1 Recommended HMC Level	Recommended HMC and system firmware combination for the SF240 Release Level.	Supported HMC and system firmware combination. There are currently no planned releases of Service Packs for this Release Level (SF235). Recommend upgrading to Release Level SF240.	Supported HMC and system firmware combination. There are currently no planned releases of Service Packs for this Release Level (SF235). Recommend upgrading to Release Level SF240.	Recommend upgrading to Release Level SF240.
HMC V5R2 Minimum HMC Level required to support POWER5 Release Level SF240. Recommend upgrading HMC Level to V6R1	Recommend upgrading HMC Level to V6R1.	There are currently no planned releases of Service Packs for this Release Level (SF235). Recommend upgrading to Release Level SF240 and HMC Level V6R1.	There are currently no planned releases of Service Packs for this Release Level (SF230). Recommend upgrading to Release Level SF240 and HMC Level V6R1.	Recommend upgrading to Release Level SF240 and HMC Level V6R1.
HMC V5R1 Minimum HMC level required to support POWER5 Release Level 235. Recommend upgrading HMC Level to V6R1	Not a supported combination.	There are currently no planned releases of Service Packs for this Release Level (SF235). Recommend upgrading to Release Level SF240 and HMC Level V6R1.	There are currently no planned releases of Service Packs for this Release Level (SF230). Recommend upgrading to Release Level SF240 and HMC Level V6R1.	Recommend upgrading to Release Level SF240 and HMC Level V6R1.
HMC V4R5 Minimum HMC level required to support POWER5 Release Level 230 Recommend upgrading HMC Level to V6R1	Not a supported combination.	Not a supported combination.	There are currently no planned releases of Service Packs for this Release Level (SF235). Recommend upgrading to Release Level SF240 and HMC Level V6R1.	Recommend upgrading to Release Level SF240 and HMC Level V6R1.

Matrix Key:

Latest Release Level
Maximum Stability Release Level
Reduced Fix support
End of Service Pack support

Supported HMC and Server code combinations : For 590 and 595 only

The following table lists currently supported firmware (FW) Release Levels for POWER5 systems, as well as the compatibility of HMC FW levels with system FW levels.

HMC levels	POWER 5 system firmware levels (iSeries and pSeries)			
	240 Release	235 Release	230 Release	225 Release
HMC V7R3 Minimum HMC level to support POWER6. Also supports POWER5 systems at Release Level SF240_299 and above	Supported combination for Service Pack SF240_299 and higher	Not a recommended or supported combination	Not a recommended or supported combination	Not a recommended or supported combination
HMC V6R1 Recommended HMC Level	Recommended HMC and system firmware combination for the SF240 Release Level.	Supported for 595 or 590 servers shipped from IBM with the FW 235 release installed. There are currently no planned releases of Service Packs for this Release Level (SF235). Recommend upgrading to Release Level SF240.	Supported HMC and system firmware combination. There are currently no planned releases of Service Packs for this Release Level (SF235). Recommend upgrading to Release Level SF240.	Recommend upgrading to Release Level SF240.
HMC V5R2 Supported HMC level for SF235 and SF230. Minimum HMC Level required to support POWER5 Release Level SF240_202 and SF240_219. Recommend upgrading HMC Level to V6R1	Recommend upgrading HMC Level to V6R1.	There are currently no planned releases of Service Packs for this Release Level (SF235). Recommend upgrading to Release Level SF240.	There are currently no planned releases of Service Packs for this Release Level (SF230). Recommend upgrading to Release Level SF240 and HMC Level V6R1.	Recommend upgrading to Release Level SF240 and HMC Level V6R1.
HMC V5R1 Minimum HMC level required to support POWER5 Release Level 235. Recommend upgrading HMC Level to V6R1	Not a supported combination.	There are currently no planned releases of Service Packs for this Release Level (SF235). Recommend upgrading to Release Level SF240.	There are currently no planned releases of Service Packs for this Release Level (SF230). Recommend upgrading to Release Level SF240 and HMC Level V6R1.	Recommend upgrading to Release Level SF240 and HMC Level V6R1.
HMC V4R5 Minimum HMC level required to support POWER5 Release Level 230. Recommend upgrading HMC Level to V6R1	Not a supported combination.	Not a supported combination.	There are currently no planned releases of Service Packs for this Release Level (SF235). Recommend upgrading to Release Level SF240.	There are currently no planned releases of Service Packs for this Release Level (SF230). Recommend upgrading to Release Level SF240 and HMC Level V6R1.

Matrix Key:

	Latest Release Level
	Maximum Stability Release Level
	Reduced Fix support
	End of Service Pack support

59x Upgrade Path Reference Table

59x Upgrade Path Reference Table		
From Level	To Level	Mechanism
SF230_XXX	SF235_XXX	Not supported
SF230_XXX	SF240_XXX	Order ECA815 and an IBM SSR will perform the upgrade (as required)
SF235_XXX	SF240_XXX	9119 -59x Systems - Order ECA834 and an IBM SSR will perform the upgrade (mandatory)
SF235_XXX	SF240_XXX	9406-595 Systems - Customer upgradeable or order ECA834 and an IBM SSR will perform the upgrade (as required)

IBM introduced the Concurrent Firmware Maintenance (CFM) function on p5 systems in system firmware level SF230_126_120, which was released on June 16, 2005. This function supports non disruptive system firmware service packs to be applied to the system concurrently (without requiring an IPL to activate changes).

For systems that are not managed by an HMC, the installation of system firmware is always disruptive.

Note: The concurrent levels of system firmware may, on occasion, contain fixes that are known as deferred. These deferred fixes can be installed concurrently, but will not be activated until the next IPL. Deferred fixes, if any, will be identified in the "Firmware Update Descriptions" table of this document. For deferred fixes within a service pack, only the fixes in the service pack which cannot be concurrently activated are deferred.

Use the following example as a reference to determine whether your installation will be concurrent or disruptive.

Note: The file names and service pack levels used in the following examples are for clarification only, and are not necessarily levels that have been, or will be released.

System firmware file naming convention for Power 5 machines : 01SFXXX_YYY_ZZZ

XXX is the release level
 YYY is the service pack level
 ZZZ is the last disruptive service pack level

NOTE: Values of service pack and last disruptive service pack level (YYY and ZZZ) are only unique within a release level (XXX).

System firmware file naming convention for Power 6 machines : 01EMXXX_YYY_ZZZ

XXX is the release level
 YYY is the service pack level
 ZZZ is the last disruptive service pack level.

NOTE: Values of service pack and last disruptive service pack level (YYY and ZZZ) are only unique within a release level (XXX).

An installation is **disruptive** if:

- The release levels (XXX) are different.
- The service pack level (YYY) and the last disruptive service pack level (ZZZ) are equal.
- The service pack level (YYY) currently installed on the system is lower than the last disruptive service pack level (ZZZ) of the service pack to be installed.

An installation is **concurrent** if:

- The service pack level (YYY) is higher than the service pack level currently installed on your system.

IO Firmware Upgrade.

It is comfortable in upgrading the IO Frimware using the “I/O Microcode Update Files and Discovery Tool CD-ROM image”.

The CD-ROM based MDS is designed for systems which are not internet-connected. Like the Web based version, it is used to determine if your IBM System p or RS/6000 is at the latest microcode level.

Mount the CD-ROM by entering the following command: *mount -rv cdrfs /dev/cd0 /cdrom*

Run the Microcode Update Tool from the CD using the following commands.
Go to the direcotry */tmp/microcode/RPM* and clear the files if present inside.
Then cd to */cdrom* or */mnt* where ever the CD is mounted.

./MDT -x

The Microcode Update Tool performs the following tasks:

Surveys the microcode levels installed on your system
Compares these levels with the latest levels available from the MDT CD-ROM.

```
bash-2.05b# ./MDT -x
DISPLAY=bgl-svr-ebs-02:0.0
CD-ROM Microcode Survey and Copy Tool:
command-line mode
CD Home: '/ml/OS/FW/pSeries_FW_Jul_18_2007/IOFW/.'
Microcode catalog date : 2007-07-10
installed Inventory Scout version=2.2.0.10
installed Logic Data Base version=2.2.0.2
Microcode catalog on host: 2007-07-10
Aix version=5.2
Inventory Scout does not need to be updated on this host.
Logic DB is current on this host.

Running InvScout:
InvScout run finished.
Analyzing survey file '/var/adm/invscout/bgl-svr-ebs-02.mup'

No microcode packages are out of date on this host.
Enter: L=List all, X=eXit : L

package : 7028-6C4;7028-6E4 [7028-6C4-system]
  readme : /ml/OS/FW/pSeries_FW_Jul_18_2007/IOFW/./microcode/M7028-6C4/70286C4F.txt
  device : system
installed : 3R070425
  latest : 3R070425
  action : None

package : PCI-X Dual Channel Ultra320 SCSI Adapter [n/a]
  readme : /ml/OS/FW/pSeries_FW_Jul_18_2007/IOFW/./microcode/F5702_A/sisscsia070800
  device : sisscsia0
installed : 0708000b
  latest : 0708000b
  action : None
2 of 9 : Q=Quit, <return> for more █
```

Retrieves the microcode from the Microcode Update files and Discovery Tool CD-ROM.

Inventory Scout (invscout) is a low-level utility, distributed with AIX, which does the actual polling of devices to determine their microcode levels. Following are the suggested actions from the MDT tool.

Update	Indicates the installed version is downlevel and needs to be updated.
None	Indicates the installed version is the latest available and no action is required.
Inspect	The installed version is different from the latest available version in our configuration file. In most cases, this difference means the installed version is downlevel and needs to be updated. But there is a small possibility that the installed level is actually higher than the latest available level in our configuration file. In this case, you do not need to update. One hint is the "Release Date". If you know (from another source) that the installed version has a later "release date" than the release date listed in the Results page, the installed version is probably up to date. View the README for the individual microcode for more detailed information on determining if it needs to be installed.
Research	The installed version cannot be determined by the Microcode Discovery Service. Follow the instructions in the README file to determine the installed version.

Use the options in the MDT itself to copy the microcode rpm files.

If the update tool suggest the upgrade method as "Research", please do avoid that microcode.

You can attempt to run the MDT tool again after inventory scout program upgrade and one round of system microcode and IO microcodes upgrade, it may make microcodes with type research to be detected as a suggested upgrade.

Go to the directory */tmp/microcode/RPM*.

After you have copied the microcode to disk, review each ReadMe file to understand and complete the microcode installation process.

Generally the following method helps to install the microcodes.

Install the microcode rpms : *rpm -ivh -ignoreos *.rpm*.

Then use the *diag* tool and use the *Microcode Download tasks* to update the microcodes.

FUNCTION SELECTION

801002

Move cursor to selection, then press Enter.

Diagnostic Routines

This selection will test the machine hardware. Wrap plugs and other advanced functions will not be used.

Advanced Diagnostics Routines

This selection will test the machine hardware. Wrap plugs and other advanced functions will be used.

Task Selection (Diagnostics, Advanced Diagnostics, Service Aids, etc.)

This selection will list the tasks supported by these procedures. Once a task is selected, a resource menu may be presented showing all resources supported by the task.

Resource Selection

This selection will list the resources in the system that are supported by these procedures. Once a resource is selected, a task menu will be presented showing all tasks that can be run on the resource(s).

F1=Help

F10=Exit

F3=Previous Menu

TASKS SELECTION LIST

801004

From the list below, select a task by moving the cursor to the task and pressing 'Enter'.

To list the resources for the task highlighted, press 'List'.

[MORE...30]

This selection provides a set of tools to perform utility functions on Tape and Medium Changer Devices supported by the Atape Device Driver.

Identify and Attention Indicators

Local Area Network Analyzer

Log Repair Action

Microcode Tasks

RAID Array Manager

SCSI BUS Analyzer

SSA Service Aids

This selection provides tools for diagnosing and resolving problems on SSA attached devices.

Update Disk Based Diagnostics

[BOTTOM]

F1=Help

F4=List

F10=Exit

Enter

F3=Previous Menu

```

Microcode Tasks

Move cursor to desired item and press Enter.

  Display Microcode Level
  Download Latest Available Microcode
  Download Microcode
  Generic Microcode Download

```

```

RESOURCE SELECTION LIST 801006

From the list below, select any number of resources by moving
the cursor to the resource and pressing 'Enter'.
To cancel the selection, press 'Enter' again.
To list the supported tasks for the resource highlighted, press 'List'.

Once all selections have been made, press 'Commit'.
To avoid selecting a resource, press 'Previous Menu'.

[MORE...12]
  ent8          P1-T6          2-Port 10/100/1000 Base-TX PCI-X Adapter
                  (14108902)
  ent9          P1-T7          2-Port 10/100/1000 Base-TX PCI-X Adapter
                  (14108902)
                U7879.001.DQDFZYN-
+ fcs3          P1-C1-T1       FC Adapter
  sisscsia4     P1            PCI-X Ultra320 SCSI Adapter
  hdisk2        P1-T12-L4-L0   16 Bit LVD SCSI Disk Drive (73400 MB)
[MORE...28]

F1=Help          F4=List          F7=Commit        F10=Exit
F3=Previous Menu

```

```

INSTALL MICROCODE
fcs3    FC Adapter

The current microcode level for fcs3 is 190104.

Available levels to install are listed below.
Select the microcode level to be installed.

Use Help for explanations of "M", "L", "C"
and "P" .

Make selection, use Enter to continue.

M 191105

```

```
INSTALL MICROCODE
fcs3      FC Adapter

Microcode installation is in progress.

Please stand by.█
```

```
INSTALL MICROCODE
fcs3      FC Adapter

Installation of the microcode has completed successfully.
The current microcode level for fcs3 is 191105.

Please run diagnostics on the adapter to ensure that it is
functioning properly.

Use Enter to continue.█
```

You need to do it on all the lpars in a partitioned Managed Server. Unfortunately this also cannot ensure that the IO cards which are not assigned to any lpars will get upgraded.

AIX Tuning.

As you all know AIX tuning is a very vast subject. Tuning involves Operating System subsystems like CPU, Memory, Disk IO, Network IO and User environment subsystems like number of processes, ptys, command arguments etc.. Topics covered under AIX Tuning here involves

System Environment.

We are not discussing the fundamentals behind the below recommended tuning parameters. But 90-95% of the cases, the following are the suggested system tuning parameters.

Performance

Run the following commands in the below box as root.

```

###Networking Tuning Parameters####
#####
no -p -o tcp_recvspace=262144
no -p -o tcp_sendspace=262144
no -p -o udp_recvspace=655360
no -p -o udp_sendspace=65536
no -p -o rfc1323=1
#####

no -L | grep -e space -e rfc

###Virtual Memory Tuning Parameters####
#####
vmo -p -o maxperm%=90
vmo -p -o maxclient%=90
vmo -p -o minperm%=5
vmo -p -o lru_file_repage=0
vmo -p -o lru_poll_interval=10
#####

vmo -L | grep -e perm -e lru -e client

```

Name Resolution.

Edit the file */etc/netsvc.conf* and add the following line to it.

hosts = local, bind

```
hosts = local, bind
```

If DNS servers need to be configured, ensure that the */etc/resolv.conf* has correct valid entries.

User Environment.

i) Edit */etc/profile* :

Edit */etc/profile* and include the following new lines :

```
stty erase ^?
set -o vi
TMOU=600 ; TIMEOUT=600 ; export readonly TMOU TIMEOUT
```

ii) Edit a users *.profile* file:

Edit a users *.profile* and include the following new lines :

```
export HOSTNAME=`hostname`
export PS1='${USER}@${HOSTNAME}:${PWD} $ '
```

iii) Maximum number of PROCESSES allowed per user

The default maxuproc value of 128 may not be sufficient enough for some environment.
It is a good practice to increase it to 1024 and if required to 4096.

```
chdev -l sys0 -a maxuproc=1024
```

iv) ARG/ENV list size

The default ncargs value of 6 may irritate some administrators.
Can change it to a higher value.

```
chdev -l sys0 -a ncargs=32
```

v) Setting user resource limits, the ulimit.

Increasing system from its default limits in line with the system resources is always a good practice.
We had seen even a file backup fails due to the reason that the default limits are too small.
Keeping limits as unlimited is also fine but some times it is potentially dangerous.

Unless otherwise specified, the following limits are good enough in a 8GB RAM machine.

```
$ ulimit -a
time(seconds)      unlimited
file(blocks)       8388608
data(kbytes)       6291456
stack(kbytes)      1048576
memory(kbytes)     4194304
coredump(blocks)   2097151
nofiles(descriptors) 20000
```

Check your current ulimit with the command *ulimit -a*,
for Soft limit *ulimit -Sa* & for Hard limit *ulimit -Ha*

Driver Installation / Upgrades

AIX is a comprehensive Operating System. Almost all hardware devices supported by AIX have their drivers already present in the system by default. Else a *cfgmgr -i <AIX lpps>* is enough. But there are 2 instances where we need to install separate AIX drivers. For external Storage Subsystems including IBM System Storage DS6000/ DS8000 / ESS / SVC as well as from other vendors and external Tape library including IBM Tape Storage. IBM System Storage DS4000 Series (earlier known as FASTT) drivers are present in the AIX by default. Installation of these drivers can be done using “*smitty installp*” utility.

Fortunately for the entire IBM Ultrium Tape Storage Libraries only a single driver package (Atape) is enough. Atape is a very powerful and flexible driver and utility stack, capable enough to manage the entire library as an Enterprise backup system when used along with UNIX utilities and a little bit UNIX scripting.

The IBM Ultrium tape and medium changer device drivers are designed specifically to take advantage of the features provided by the IBM Ultrium tape drives and medium changer devices. The goal is to give applications access to the functions required for basic tape functions (such as backup and restore) and medium changer operations (such as cartridge mount and demount), as well as to the advanced functions needed by full tape management systems. Whenever possible, the driver is designed to take advantage of the device features transparent to the application. Atape driver supports the following IBM Library products.

- IBM TotalStorage Ultrium External Tape Drive 3580.
- IBM TotalStorage Ultrium Tape Autoloader 3581.
- IBM TotalStorage Ultrium Tape Library 3582.
- IBM TotalStorage Ultrium Scalable Tape Library 3583.
- IBM TotalStorage UltraScalable Tape Library 3584.
- IBM Virtualization Engine TS7510.
- IBM System Storage TS3310 Tape Library.
- IBM System Storage TS3100 Tape Library.

Following are the driver packages and their versions at the time of this writing.

Driver support for Ultrium Tape Drive & Tape Libraries.

Atape driver package	–	Atape.10.5.4.0.bin	Version 10.5.4.0
----------------------	---	--------------------	------------------

IBM System Storage DS6800 / DS8000 / SVC are supported by FCP driver packages and SDD (Multipath Subsystem Device Driver) driver packages. Following are the driver packages and their versions at the time of this writing. Please do remember that the FCP and SDD driver versions are AIX Version and Technology Level dependent. Hence select the right combinations for the entire stack to work flawlessly. Moreover fcp driver should be installed or upgraded prior to the sdd driver installation or upgrade.

Driver support for DS6800 / DS8000 / SVC

FCP driver package	–	devices.fcp.disk.ibm.rte	V1.0.0.9
SDD driver package	-	devices.sdd.52.rte	V1.6.3.0
	-	devices.sdd.53.rte	V1.6.3.0

"devices.fcp.disk.ibm.rte" is a host attachment package for AIX. This package provides FCP device support for DS8000, DS6000, and SVC devices.

SDD is the IBM Multipath Subsystem Device Driver. SDD provides multipath configuration environment support for a host system that is attached to storage devices. It provides enhanced data availability, dynamic input/output (I/O) load balancing across multiple paths, and automatic path failover protection. Because SDD can be installed in so many different environments / configurations, detailed information about each environment is placed in the appropriate chapter of the Multipath SDD Users Guide. The 'Summary of Changes' section of the SDD Users' Guide can help you quickly determine if the latest changes affect you.

Starting from SDD 1.6.0.0, SDD supports the coexistence of all supported storage devices

Starting from SDD 1.6.2.0, a unique ID attribute is added to SDD vpath devices, in order to support AIX5.3 VIO future features. AIX device configure methods have been changed in both AIX52 TL08 (52M) and AIX53 TL04 (53E) for this support. Following are the requirements for this version of SDD on AIX5.2 and AIX5.3:

AIX52 TL08 & above with PTF U804193 (IY76991)
AIX53 TL04 & above with PTF U804397 (IY76997)

Please do check the following unsupported configurations with SDD.

Unsupported environments SDD does not support :

- A host system with both a SCSI and fibre-channel connection to a shared ESS LUN.
- Placing system primary paging devices (for example, /dev/hd6) on an SDD vpath device.
- Any application that depends on a SCSI-2 reserve and release device on AIX.
- Single-path mode during concurrent download of licensed machine code nor during any disk storage systems concurrent maintenance that impacts the path attachment, such as a disk storage systems host-bay-adapter replacement.
- Multipathing to a system boot device.
- Configuring SDD vpath devices as system primary or secondary dump devices.
- More than 600 SDD vpath devices if the host system is running AIX 4.3.3 or AIX 5.1.0.
- More than 1200 SDD vpath devices if the host system is running AIX 5.2 or AIX 5.3 v DS8000, DS6000, and SAN Volume Controller with SCSI connectivity.
- Multiple AIX servers without SDD-supported clustering software, such as HACMP, installed.

AIX data collection : (snap & perfpmr)

Gathers system configuration information : “snap” Command

The snap command gathers system configuration information and compresses the information into a pax file. The file can then be downloaded to disk or tape, or transmitted to a remote system. The information gathered with the snap command may be required to identify and resolve system problems. This is a built in tool in AIX base operating system.

Note: Root user authority is required to execute the snap command.

- Can compresses the information into a pax file.
- Root user authority is required to execute the snap command.
- Approximately 8MB of temporary disk space is required.
- /tmp/ibmsupt directory, the default directory for snap command output.
- You can write the output to another directory by using the -d flag.
- Each execution of the snap command appends information to previously created files.
- Use the -r flag to remove previously gathered and saved information.

Common option : snap -g

gathers general system information, including the following:

- * Error report
- * Copy of the customized (ODM) database
- * Trace file
- * User environment
- * Amount of physical memory and paging space
- * Device and attribute information
- * Security user information

The output of the snap -g command is written to the /tmp/ibmsupt/general/general.snap file.

Other snap options.

- a Gathers all system configuration information.
This option requires approximately 8MB of disk space.
- c Creates a compressed pax image (snap.pax.Z file) of all files.
- D Gathers dump and /unix information. The primary dump device is used.
- d <AbsolutePath> Identifies the optional snap command output directory.
- e Gathers HACMP(TM) specific information from all nodes.
- r Removes snap command output from the /tmp/ibmsupt directory.

Suggested snap syntax : *snap -ac -d <full-path>*

```

9.184.225.165 - PuTTY
bash-3.00# clear
bash-3.00# snap -ac -d /data/mysnap
Destination directory set to /data/mysnap
Checking space requirement for general information.....
..... done.
Checking space requirement for tcpip information.....
..... done.
Checking space requirement for nfs information..... done.
Checking space requirement for kernel information..... done.
Checking space requirement for printer information.... done.
Checking space requirement for dump information..... done.
Checking space requirement for sna information.../var/sna not found
done.
Checking space requirement for filesys information.....
..... done.
Checking space requirement for async information..... done.
Checking space requirement for lang information..... done.
Checking space requirement for XS25 information.....

```

```

9.184.225.165 - PuTTY
Gathering install system information.... done.
Gathering ssa system information..... done.
Gathering logical volume manager information.....
Gathering Enhanced CLVM information.../tmp/ch.log.* not found
done.
Gathering multicpu trace filesdone.
..done.
Gathering platform/scanout information.done.
Gathering svCollect data
The script svCollect is not executable in /usr/lib/ras/snapscripts
Gathering client_collect data
Gathering lsvirt data
The script lsvirt is not executable in /usr/lib/ras/snapscripts
Gathering getRtasHeap data
Gathering pcixscsi system information..... done.

Creating compressed pax file...
Starting pax/compress process... Please wait... done.

-rw-----  1 0      0      13905415 Oct 25 06:40 snap.pax.Z
bash-3.00# █

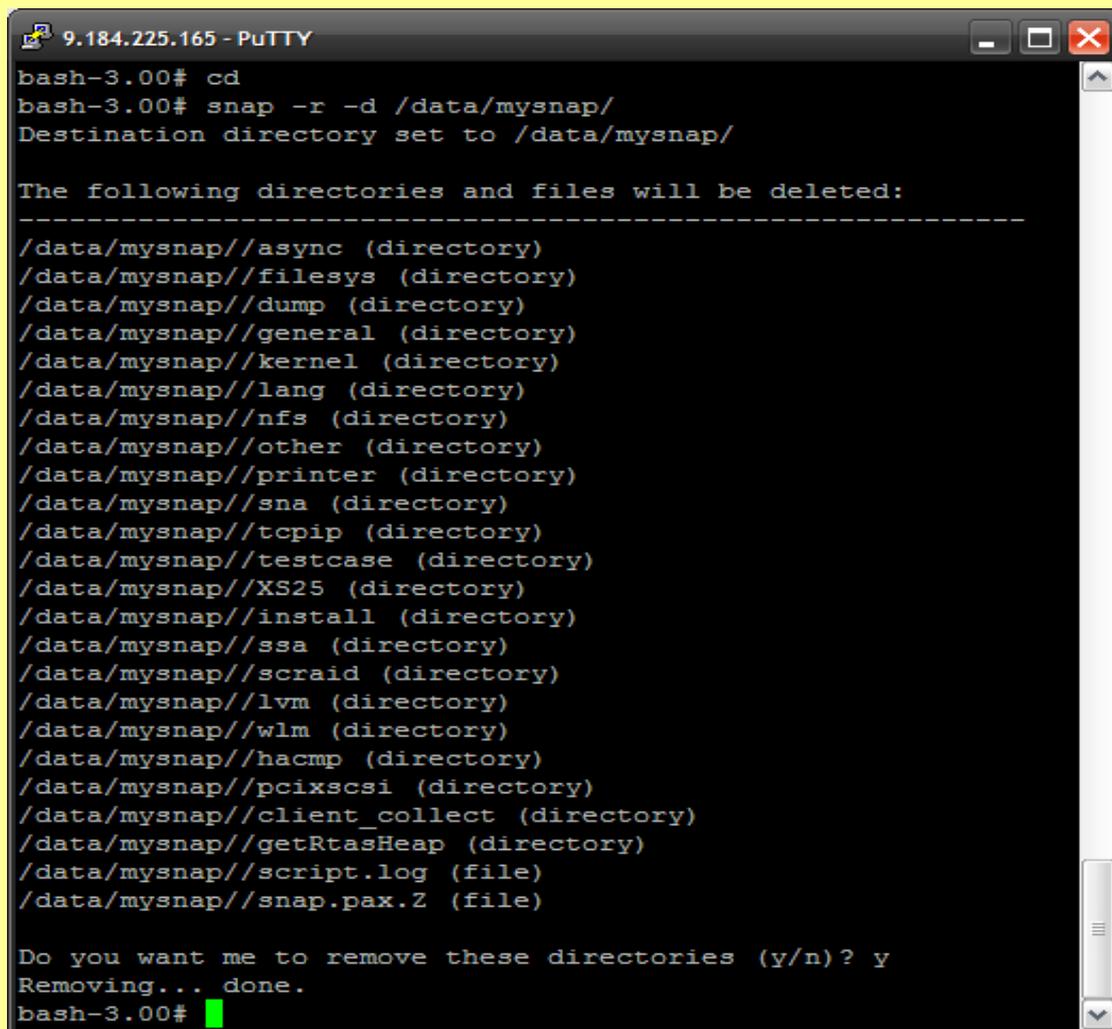
```

Together with the snap data you collected please don't forget to send the following information.

Create a file README.PROBLEM that

- a) describes the problem in detail
 - What happens?
 - What do you expect to happen?
 - What are the steps to recreate the problem?
- b) lists customer & SE/support contacts with phone numbers

You can use `snap -r` or `snap -r -d <directory>` to clear the collected snap data.



```
9.184.225.165 - PuTTY
bash-3.00# cd
bash-3.00# snap -r -d /data/mysnap/
Destination directory set to /data/mysnap/

The following directories and files will be deleted:
-----
/data/mysnap//async (directory)
/data/mysnap//filesystem (directory)
/data/mysnap//dump (directory)
/data/mysnap//general (directory)
/data/mysnap//kernel (directory)
/data/mysnap//lang (directory)
/data/mysnap//nfs (directory)
/data/mysnap//other (directory)
/data/mysnap//printer (directory)
/data/mysnap//sna (directory)
/data/mysnap//tcpip (directory)
/data/mysnap//testcase (directory)
/data/mysnap//XS25 (directory)
/data/mysnap//install (directory)
/data/mysnap//ssa (directory)
/data/mysnap//scraid (directory)
/data/mysnap//lvm (directory)
/data/mysnap//wlm (directory)
/data/mysnap//hacmp (directory)
/data/mysnap//pcixscsi (directory)
/data/mysnap//client_collect (directory)
/data/mysnap//getRtasHeap (directory)
/data/mysnap//script.log (file)
/data/mysnap//snap.pax.Z (file)

Do you want me to remove these directories (y/n)? y
Removing... done.
bash-3.00#
```

PERFORMANCE DATA COLLECTION : “perfpmr” Utility

First of all perfpmr utility is not a part of the AIX base operating system. This package contains a set of tools and instructions for collecting the data needed to analyze a AIX performance problem. They are different packages for AIX 5.2 and AIX 5.3. Always try to use the latest perfpmr package for data collection. The package will be distributed as a compressed "tar" file available electronically.

From the internet : ['ftp://ftp.software.ibm.com/aix/tools/perftools/perfpmr'](ftp://ftp.software.ibm.com/aix/tools/perftools/perfpmr)

INSTALLING THE PACKAGE

The following assumes the tar file is in /tmp and named 'perf52.tar.Z'.

- a. login as root or use the 'su' command to obtain root authority
- b. create perf52 directory and move to that directory (this ex assumes the directory built as /tmp)
`# mkdir /tmp/perf52; cd /tmp/perf52`
- c. extract the shell scripts out of the compressed tar file:
`# zcat /tmp/perf52.tar.Z | tar -xvf -`
- d. install the shell scripts : `# sh ./Install`

Collecting System Performance Data:

Detailed performance data is required to analyze and solve a performance problem. Follow these steps to invoke the supplied shell scripts:

NOTE: You must have root user authority when executing these shell scripts.

- a. Create a data collection directory and 'cd' into this directory.
 Allow at least 12MB/processor of unused space in whatever file system is used.
 Use the 'df' command to verify the filesystem has at least 15MB.

***IMPORTANT* - DO NOT COLLECT DATA IN REMOTE FILESYSTEM SINCE IPTRACE MAY HANG**

For example using /tmp filesystem:

```
# mkdir /tmp/perfdata
# cd /tmp/perfdata
```

- b. HACMP users:

Generally recommend HACMP deadman switch interval be lengthened while performance data is being collected.

- c. Collect our 'standard' PERF52 data for 600 seconds (600 seconds = 10 minutes).
 Start the data collection while the problem is already occurring with the command:

'perfpmr.sh 600'

Answer the questions in the text file called 'PROBLEM.INFO' in the data collection directory created above. This background information about your problem helps us better understand what is going wrong.

Combine all the collected data into a single binary 'tar' file and compress it:

Put the completed PROBLEM.INFO in the same directory where the data was collected (ie. /tmp/perfdata in the following example). Change to the parent directory, and use the tar command as follows:

```
# cd /tmp/perfdata (or whatever directory used to collect the data)
# cd ..
# pax -x pax -vw perfdata | gzip -c > <NAME>.pax.gz
```

NOTE:

Since a performance problems may mask other problems, it is not uncommon to fix one issue and then collect more data to work on another issue.

Questions that help IBM diagnose the problem: Extract of the "PROBLEM.INFO" file

Can you append more detail on the simplest,repeatable example of the problem?

If not, describe the least complex example of the problem.

Is the execution of AIX commands also slow?

Is this problem a case of something that had worked previously

(ie. before a upgrade) and now does not run properly?

If so: Describe any other recent changes? ie. workload, number of users, networks, configuration, etc.

Or is this a case of a application/system/hardware that is being set up for the first time?

If so:

What performance is expected?

What is the expectation based on?

Is the slow performance intermittent?

Is there any pattern to the slow behavior?

Does it get slow, but then disappear for a while?

Does it get slow at certain times of day or relation to some specific activity?

About how long does the period of slow performance before it returns to normal?

Is it slow when the system is otherwise idle? (ie. capacity vs. elapsed time)

What is the CPU utilization when the system is idle after a period of slow performance (use 'vmstat 1')? (perhaps something is looping)

Are all applications/commands/users slow or just some?

What aspect is slow?

ie. Time to echo a character,

Elapsed time to complete the transaction,,,

Does rebooting the system make the problem disappear for a while?

(ie. a resource may be consumed but not freed back up)

If so, about how long until the problem reappears?

If client/server, can the problem be demonstrated when run locally on the server (network vs. server issue)?

Does the problem disappear when the application is run from the system console?

If client/server, from the client how long does a 'ping server_ip_address' take?

(use the server_ip_address to exclude nameserver and other variables. ie. 'ping 129.35.33.22')

If network related, please describe the network segments including bandwidth

(ie. 10mbit/sec, 9600 baud,,) and routers between the client and server.

What vendor applications are on the system and are they involved in the perf issue (ie. Oracle, SAP,,)?

What is the version/release/level of the vendor applications?

Upload the data for analysis : For IBM Customers.

IBM customers can ftp the collected data for further analysis. Use binary mode for ftp.

Ftp site : *testcase.boulder.ibm.com*
 Direcotry : *toibm/aix*
 Username: anonymous
 Password : <customer's_email-address>

Download the data for analysis : Only for IBM Service Representative.

IBM engineers can download the data for analysis as follows.

```
# ftp testcase.boulder.ibm.com
Connected to testcase.boulder.ibm.com.
220-IBM's internal systems must only be used for conducting IBM's
220-business or for purposes authorized by IBM management.
220-
220-Use is subject to audit at any time by IBM management.
220-
220-This server contains IBM Confidential information. Please read
220-/README.confidential for more information.
220-
220 testcase-blue secure FTP server ready.
Name (testcase.boulder.ibm.com:dinil): dinildas@in.ibm.com
331 Password required for dinildas@in.ibm.com.
Password:
230 Virtual user dinildas@in.ibm.com logged in.
ftp> pwd
257 "/" is current directory.
ftp> cd toibm
250 CWD command successful.
ftp> cd aix
250 CWD command successful.
ftp> bin
200 Type set to I.
ftp> hash
Hash mark printing on (1024 bytes/hash mark).
ftp> mget snap.pax_sssc18_HUL_india.Z
mget snap.pax_sssc18_HUL_india.Z? y
200 PORT command successful.
150 Opening BINARY mode data connection for snap.pax_sssc18_HUL_india.Z.
#####
#####
```

Your IBM Intranet
Username / Password.