

Date : 24/08/2007	Version : 0.1
Date : 29/11/2007	Version : 0.2
Date : 30/11/2007	Version : 0.3
Date : 03/12/2007	Version : 0.4
Date : 04/12/2007	Version : 0.5
Date : 07/12/2007	Version : 0.6
Date : 08/12/2007	Version : 0.7
Date : 09/12/2007	Version : 0.8
Date : 10/05/2008	Version : 0.9.0
In progress (24/06/2008)	Version : 0.9.8.0

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) **TIMEZONE, TIME & DATE 1.0**
- 2) **Mirroring rootvg disks 1.0**
- 3) **Recommended additional packages. 1.0**
- 4) **File system structure. 1.0**
- 5) **Using lppchk - Checking AIX Integrity 1.0**
- 6) **AIX Update Strategies: MLs, TLs & SPs. 1.0**
- 7) **System Firmware Upgrade. 0.8**
- 8) **IO Firmware Upgrade. 0.7**
- 9) **AIX Tuning - System & User Environment 0.6.2**
- 10) **Driver Installation / Upgrades. 0.9**
- 11) **AIX data collection :(snap,perfpmr & sddgetdata)... 1.1**
- 12) **Configuring ESA. Planned in 1.x**
- 13) **NIM : AIX Installation & Configuration 1.0**
- 14) **Employing alt_disk_install utility 0.8**
- 15) **rpm packet management 0.7**
- 16) **Filesystem quota management 0.8**
- 17) **Configuring multibos environment 1.0**
- 18) **dlpar concordance 1.0**
- 19) **pSeries data collection. Planned in 1.x**
- 20) **A brief hardening on AIX. Planned in 2.x**
- 21) **AIX Backup & Restore mechanism Planned in 2.x**
- 22) **AIX Trusted Computing Base Planned in 2.x**
- 23) **AIX Auditing Planned in 2.x**
- 24) **AIX WLM Configuration Planned in 2.x**
- 25) **HOST Ethernet Adapter in P6 Systems Planned in 2.x**

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 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.

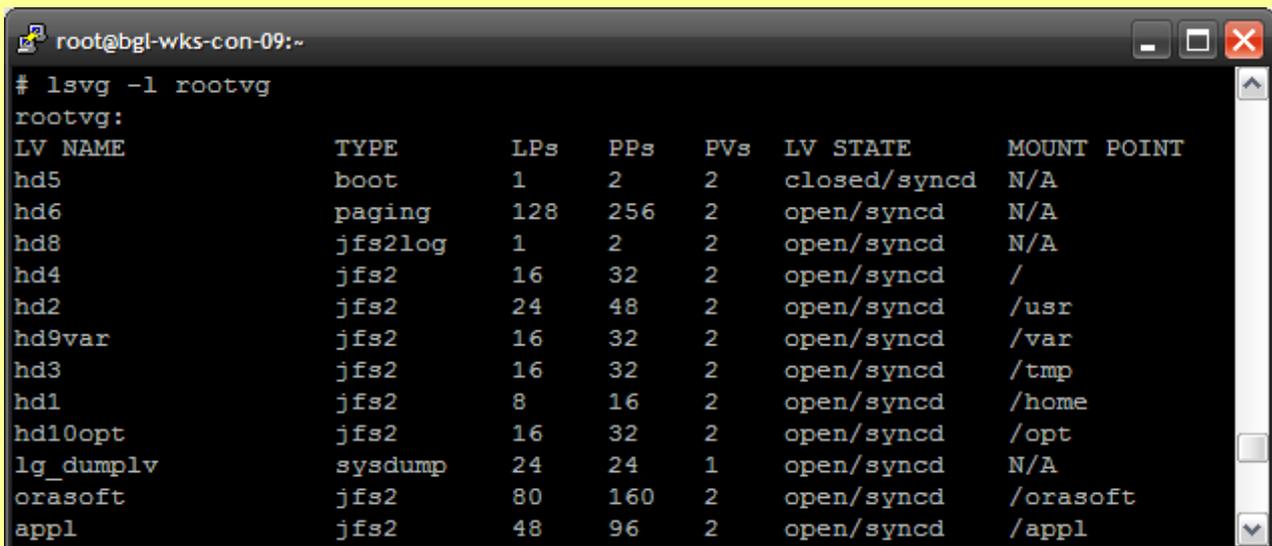
There is no need to emphasize the need for mirroring rootvg in AIX. After all who wanted to restore the mksysb backup for a simple reason of rootvg disk failure. Mirroring rootvg in AIX is a very simple & straight task. You can attempt mirroring before increasing the file system size as it takes comparatively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

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

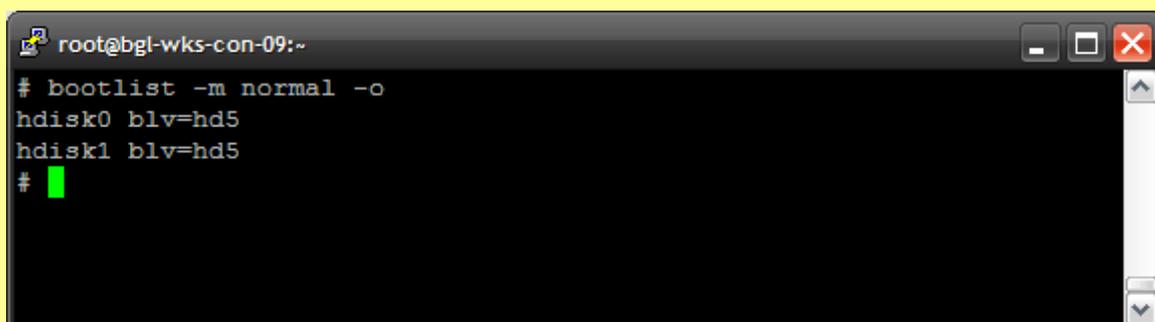
where hdiskx is the hdisk where rootvg is on and hdisky is the mirrored disk of hdiskx.

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      PPp      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_dumplv   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
#

```

Recommended additional packages.

Though the default AIX installation is comprehensive enough to address most of the functional 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.

Try to keep only one paging space per spindle.

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 system default configuration if physical memory equal to or greater than 4GB 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 system default configuration if physical memory less than 4GB 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. If there is no dedicated dump device, then on reboot, the system will attempt to copy the dump image from `/dev/hd6` to a file (`vmcore.X`) in a directory in rootvg (the default is `/var/adm/ras`). This is because the `/dev/hd6` device needs to be used as paging space when the AIX system starts. If the copy fails, often because there is not enough space, it will prompt you to copy off the dump to a tape device or to diskettes. 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 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.

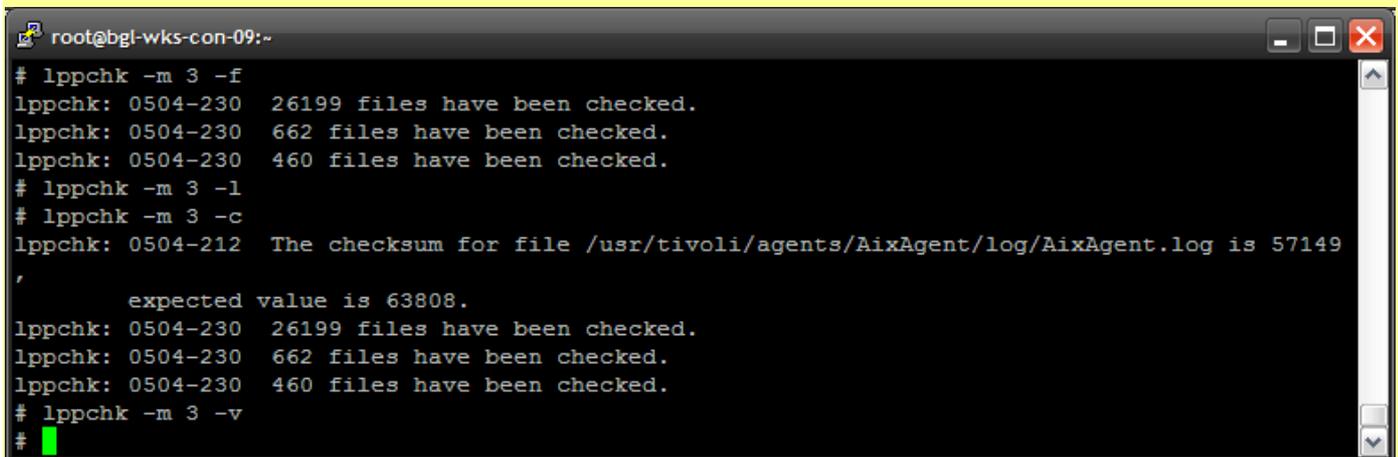
We prefer to check all the 4 types of checks for ensuring Operating System integrity.

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 trace out lack of harmonious uniformity in the system which manifests as command failures or system misbehaviors.

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:

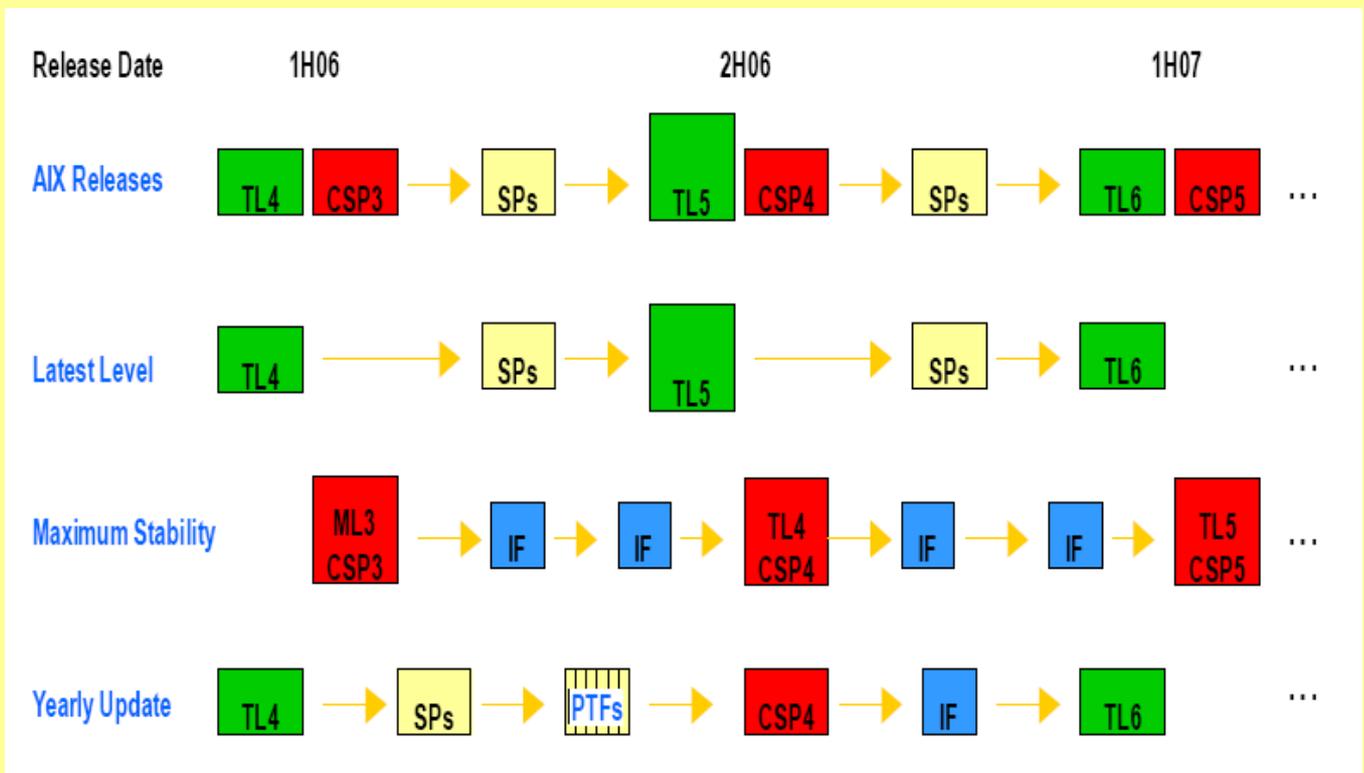
```
oslevel -s
5300-04-02
```

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:

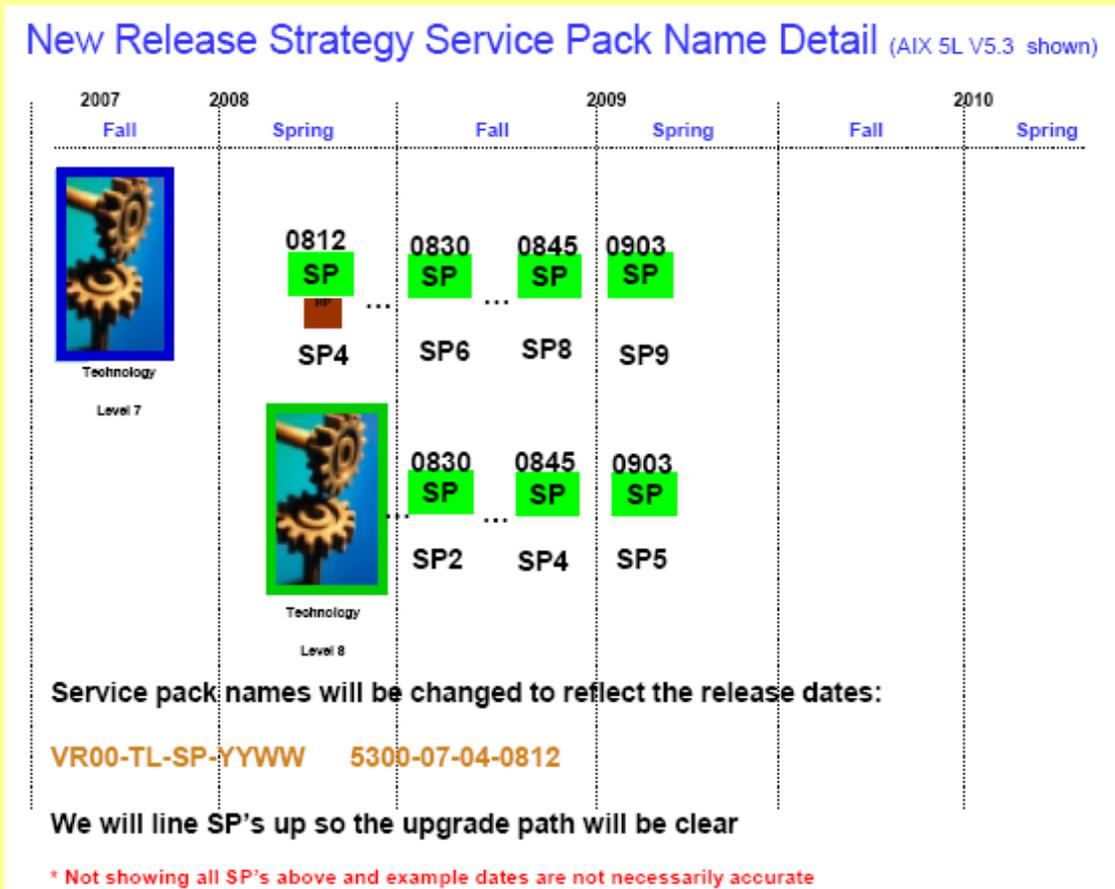
- Customer reported problems (APARs) that cannot wait until the next TL
- Critical problems found by development or test teams
- Very, very limited number of changes to support new hardware.

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 general preferable. Suggested one is the firmware which is with GA release of atleast 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_338_201** including Bulk Power **02BP240_219_168** firmware.
for **p6** machines is **EM310_069_048** and **EM320_040_031**

New Features and Functions in the P6 EM320_031_031 firmware:

Support for redundant service processors with failover on model MMA systems.

Support for the concurrent addition of a RIO/HSL adapter on model MMA systems.

Support for the concurrent replacement of a RIO/HSL adapter on model MMA systems.

Support for the creation of multiple virtual shared processor pools (VSPPs) within the one physical pool. (In order for AIX performance tools to report the correct information on systems configured with multiple shared processor pools, a minimum of AIX 5.3 TL07 or AIX 6.1 must be running.)

Support for the capability to move a running AIX or Linux partition from one system to another compatible system with a minimum of disruption.

Support for the collection of extended I/O device information (independent of the presence of an operating system) when a system is first connected to an HMC and is still in the manufacturing default state.

Improved VPD collection time on model MMA systems.

Support for EnergyScaletm and Active Energy Manager tm.

Note : Upgrade HMC code to Version 7, Release 3.2.0 before attempting to load this system firmware

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.

Generally firmware upgrade POWER6 machines takes lesser time than POWER5.

Typically a single/4 CEC 9117-MMA takes around 15-20 minutes of upgrade time.

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 the upgrade.

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 SF 240.	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 SF 240.	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 SF 240 and HMC Level V6R1.	There are currently no planned releases of Service Packs for this Release Level (SF230). Recommend upgrading to Release Level SF 240 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 SF 240 and HMC Level V6R1.	There are currently no planned releases of Service Packs for this Release Level (SF230). Recommend upgrading to Release Level SF 240 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 SF 240 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)

For Mid-Range IBM Systems with POWER6 processors

For MTMs 9117-MMA and 9406-MMA				
POWER6 system firmware levels				
HMC levels	EM320 Release		EM310 Release	
	9117-MMA	9406-MMA	9117-MMA	9406-MMA
HMC V7R3.2 + MH01084 Recommended HMC level	Supported HMC and FW combination for this model.	The EM320 FW is unavailable and not supported for this model.	Recommended and supported combination for this model, with FW level EM310_063	Recommended and supported combination for this model, with FW level EM310_063
HMC V7R3.1	Not a supported combination	The EM320 FW is unavailable and not supported for this model.	Supported combination. Recommend updating to EM310_063 FW level and HMC V7.3.2 + MH01084.	Supported combination. Recommend updating to EM310_063 FW level and HMC V7.3.2 + MH01084.

For Entry-Level IBM Systems with POWER6 processors

For MTMs 8203-E4A and 8204-E8A		
POWER6 system firmware levels		
HMC levels	EL320 Release	
	8203-E4A	8204-E8A
HMC V7R3.2 + MH01084 Recommended HMC level	Recommended and supported HMC and FW combination for this model	Recommended and supported HMC and FW combination for this model

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.

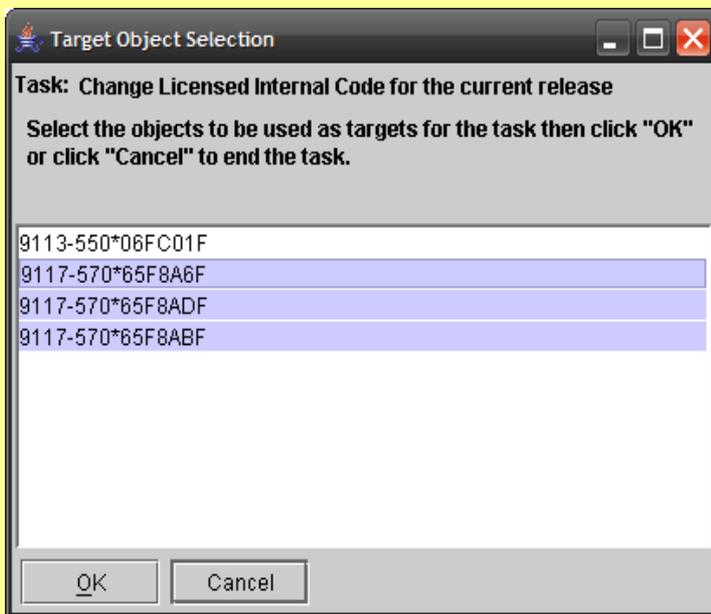
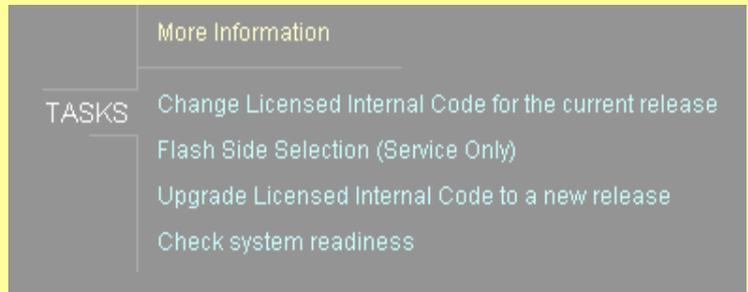
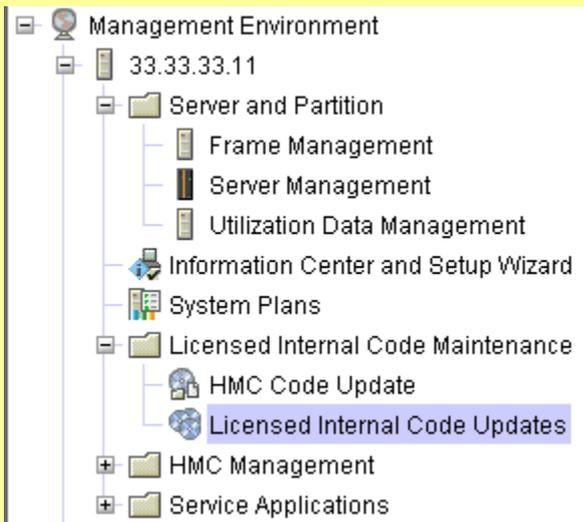
An installation is **concurrent** if:

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

POWER5 Firmware filenames and IBM internal code names

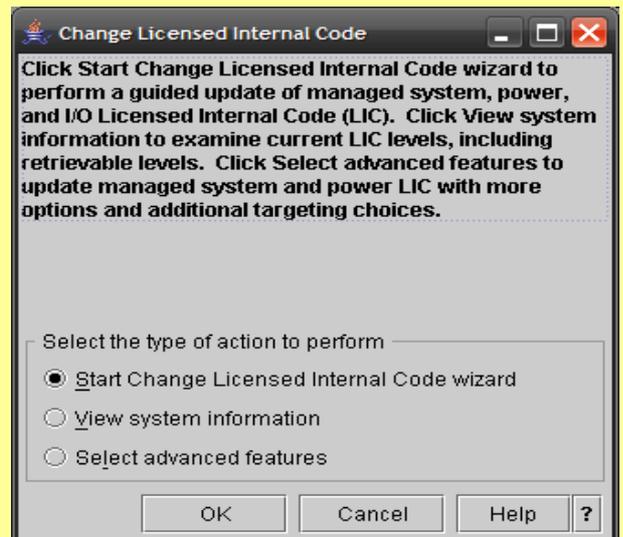
Firmware Filenames	Internal Codenames	Notes
01SF240_320_201	GA7 SP6	
01SF240_299_201	GA7 SP5+	
01SF240_298_201	GA7 SP5	
01SF240_284_201	GA7 SP4	
01SF240_261_201	GA7 SP3.2	
01SF240_259_201	GA7 SP3.1	
01SF240_258_201	GA7 SP3 (ECA 815)	
01SF240_233_201	GA7 SP2	
01SF240_222_201	GA7 SP1+	
01SF240_219_201	GA7 SP1	
01SF240_202_201	GA7+	
01SF240	GA7	
01SF235_214_160	GA6 SP5	end of SP support 3/31/07
01SF235_209_160	GA6 SP4	
01SF235_206_160	GA6 SP3 (ECA 828)	
01SF235_185_160	GA6 SP2 (ECA 821)	
01SF235_180_160	GA6 SP1	
01SF230_158_120	GA5 SP6	end of SP support 3/31/07
01SF230_156_120	GA5 SP5 (ECA 827)	
01SF230_153_120	GA5 SP 4+ (ECA 823)	
01SF230_150_120	GA5 SP 4	
01SF230_145_120	GA5 SP 3	
01SF230_143_120	GA5 SP 2	
01SF230_126_120	GA5 SP 1	
01SF230_120_120	GA5	
01SF225_096_096	GA4	

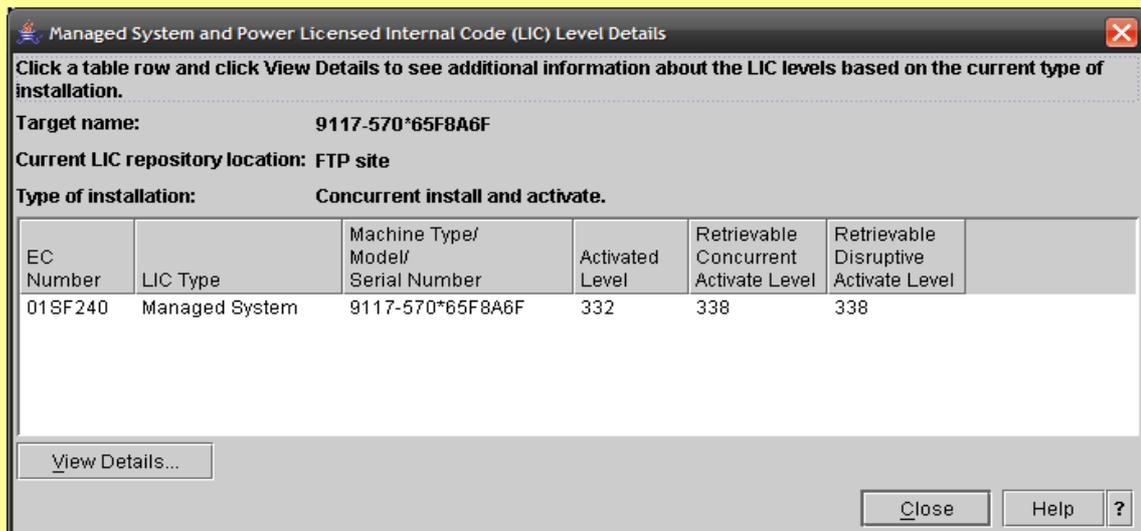
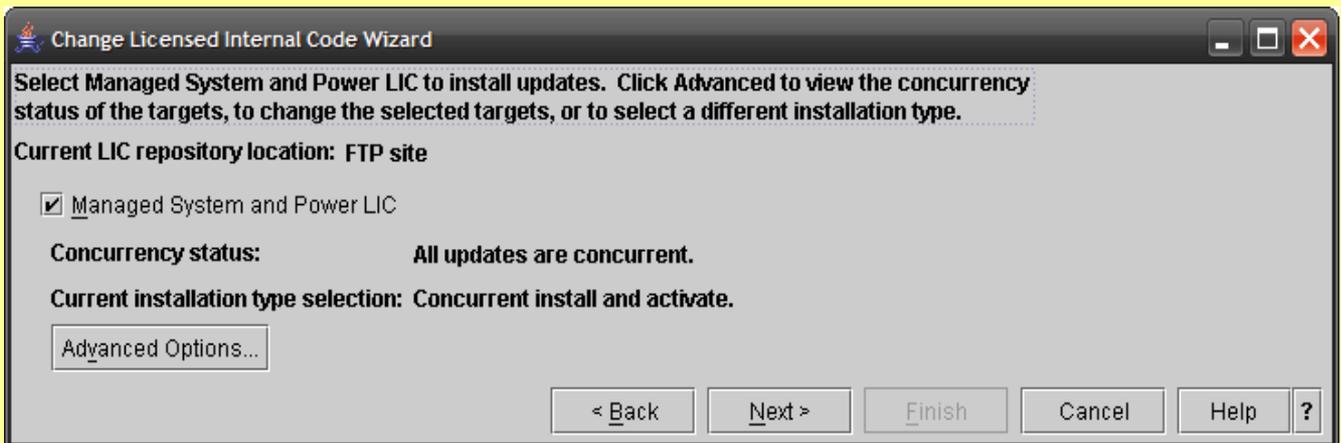
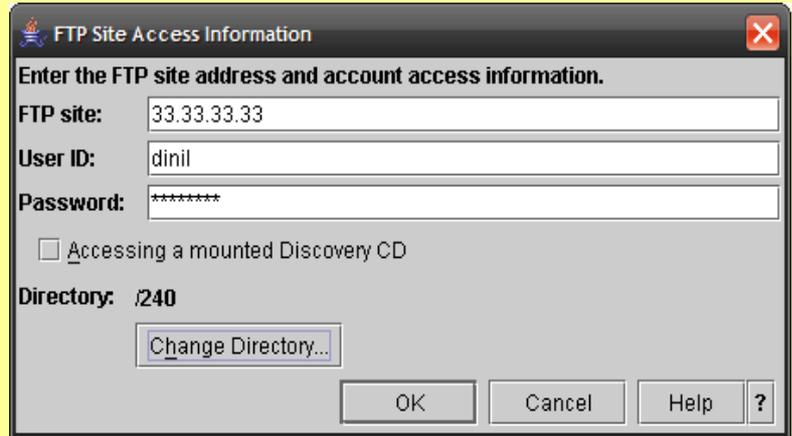
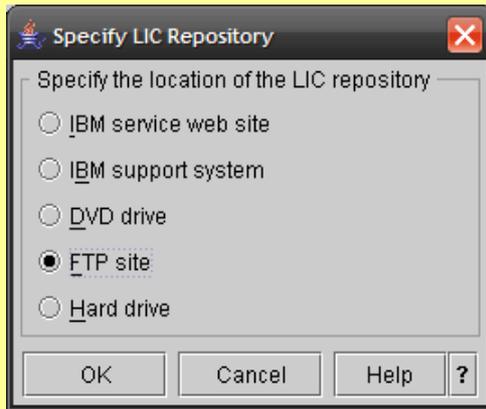
System Firmware upgrade using HMC : from a ftp LIC repository.



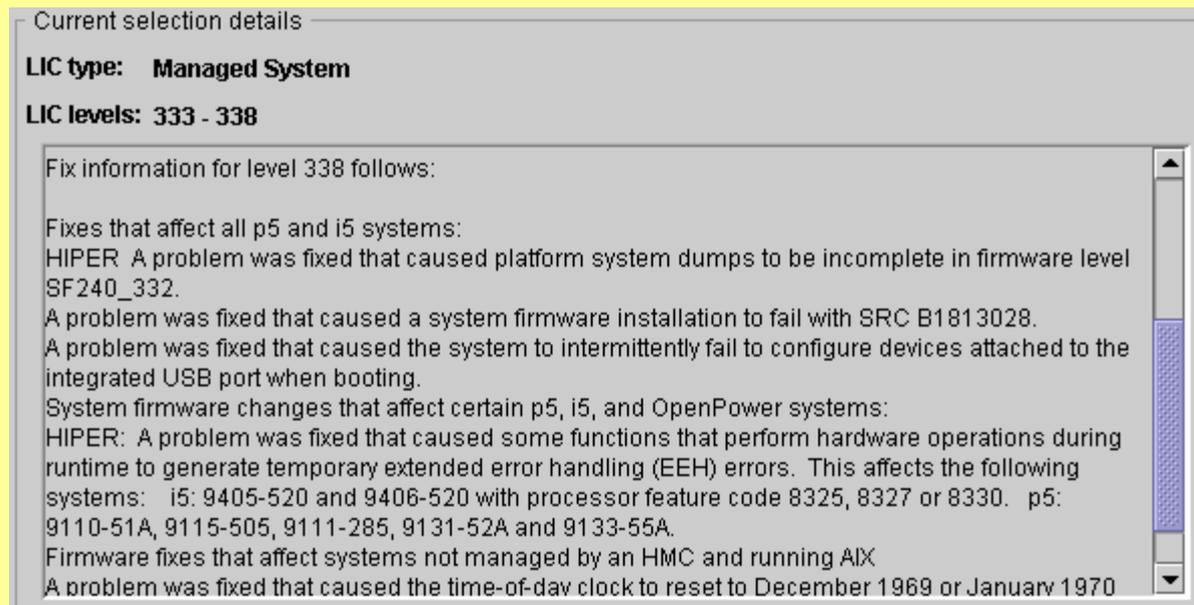
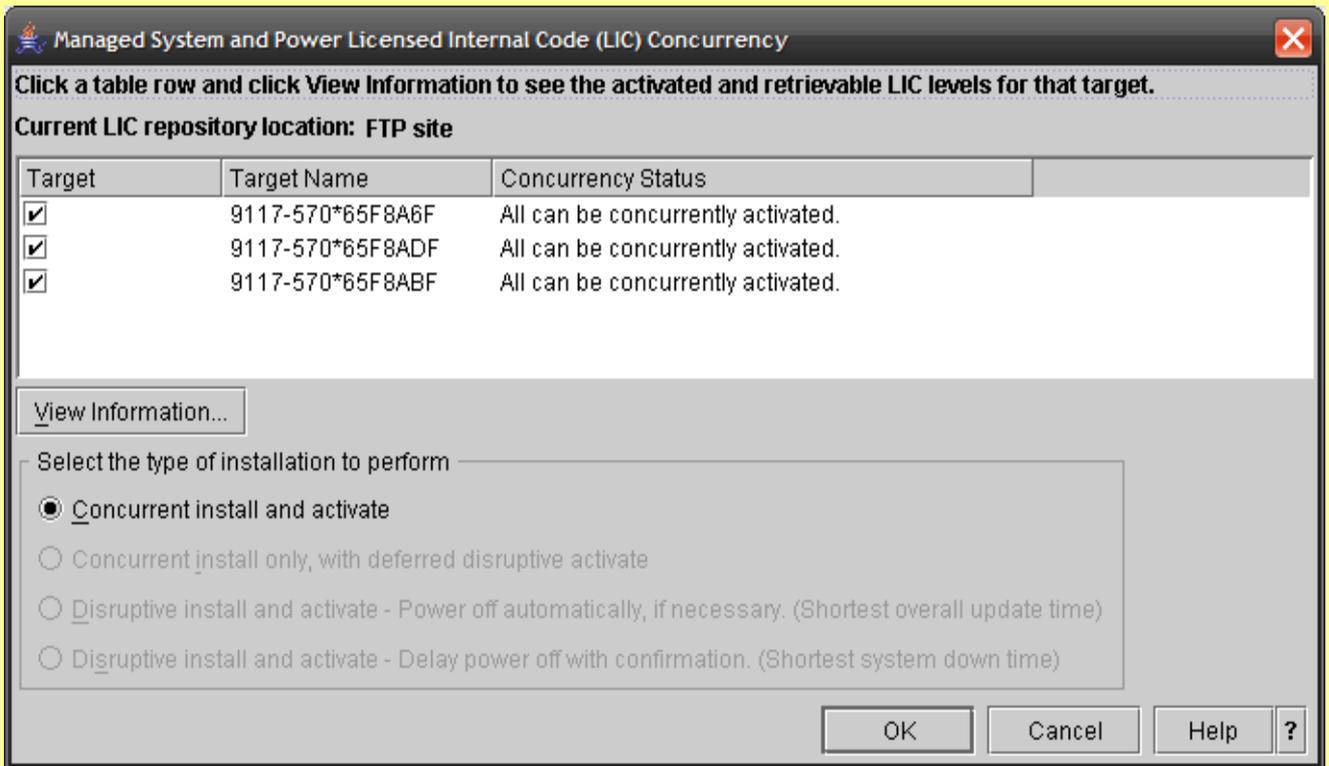
Select the 1st option for changing the LIC within the current EC level. Opt for the 3rd option for an EC level change. Remember that the Upgrade LIC option is always disruptive. Suggested to run a readiness check before attempting LIC level changes. For LIC updates as well as for readiness check, multiple managed systems can be selected together. In the depicted example we are using ftp method to update LIC to level 240_338.

In this example, a laptop is used as ftp server running FileZilla ftp service. Just a little bit of the ftp server config, with a user having a home directory with a folder named 240 where both the .rpm and .xml files are kept.

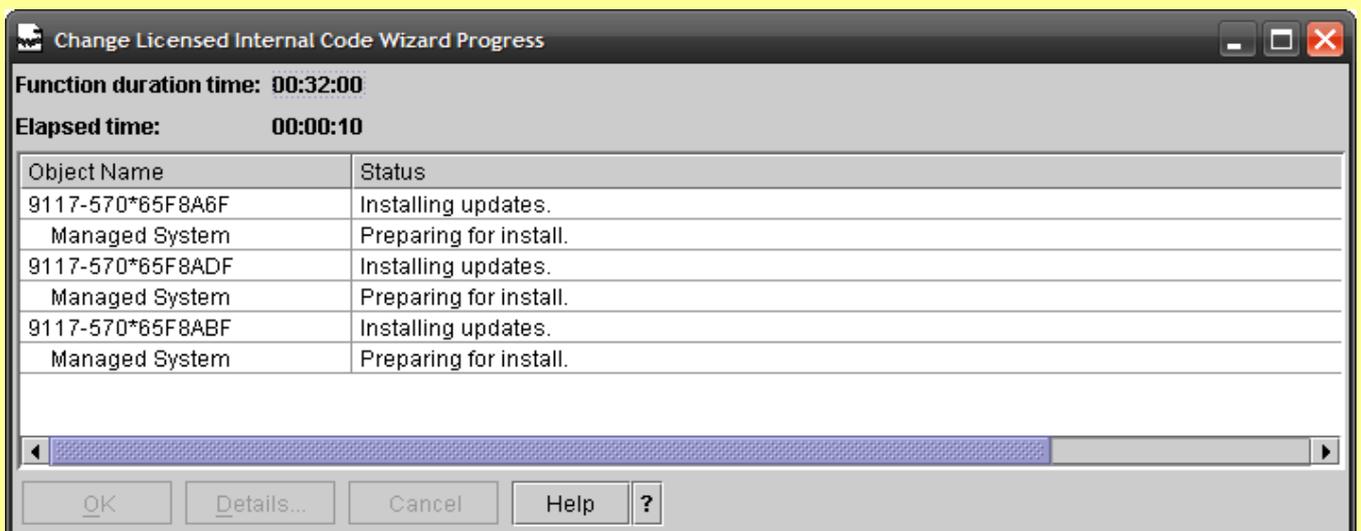
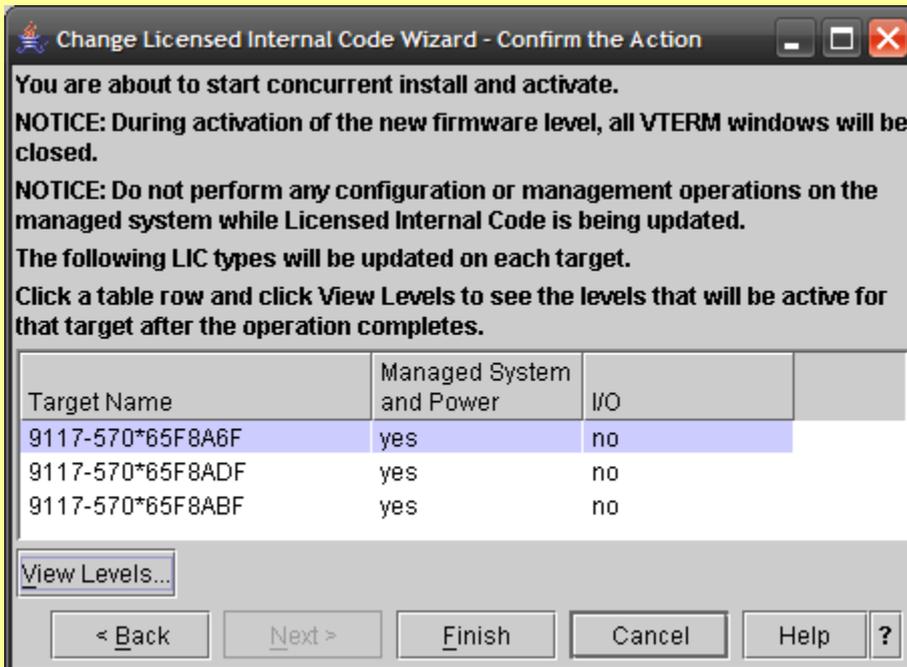
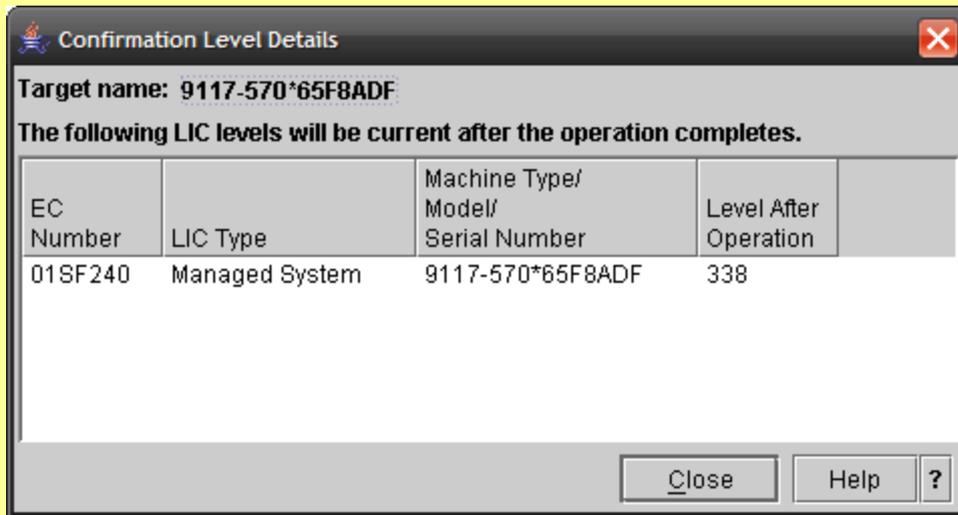


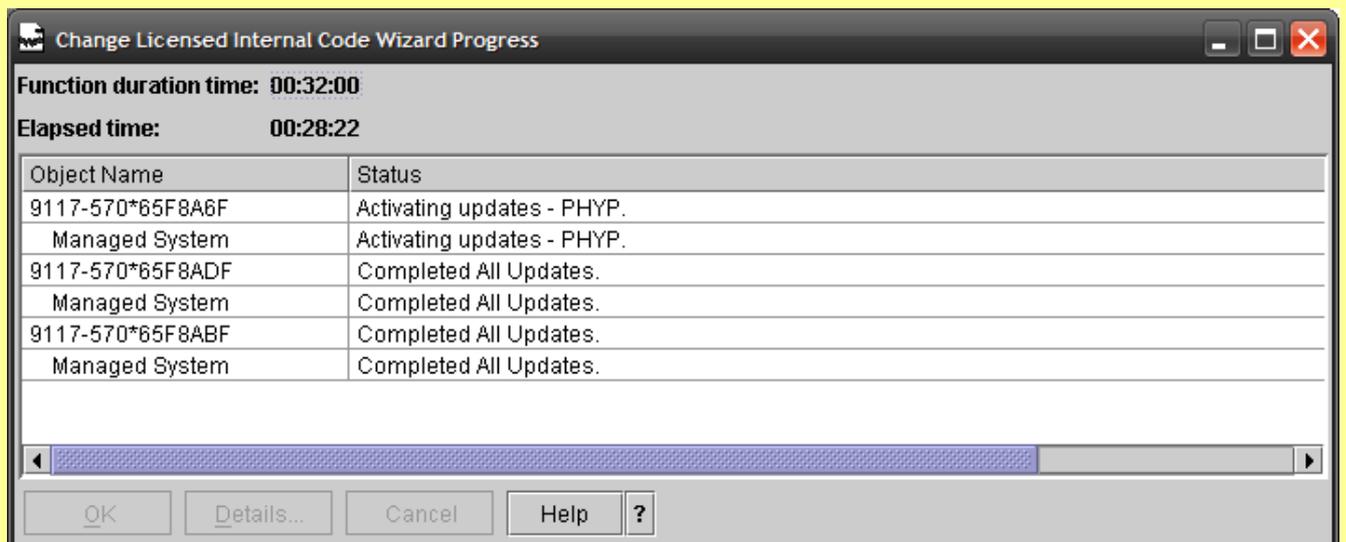
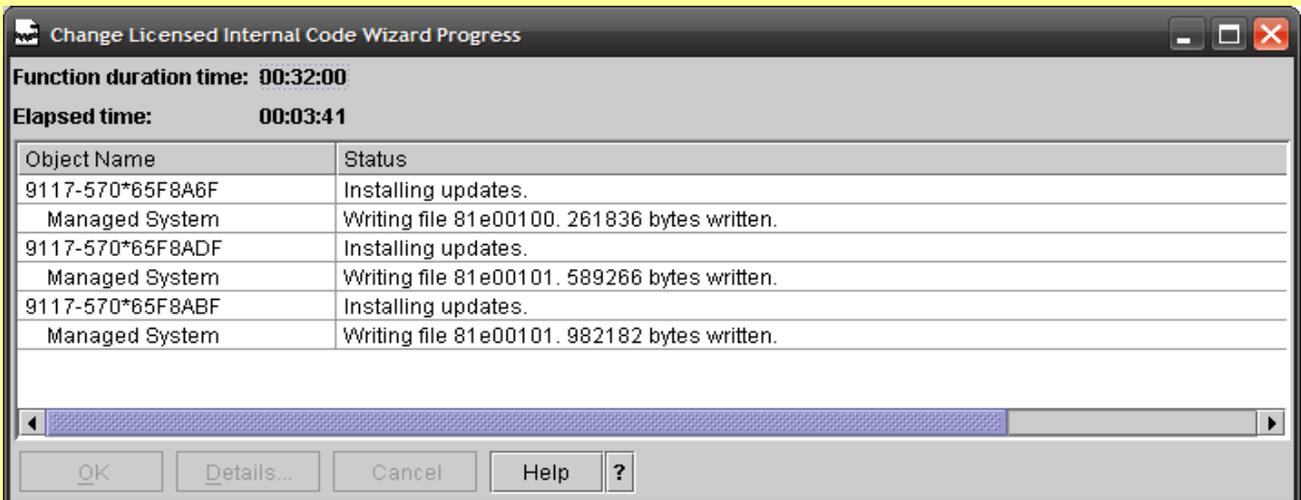
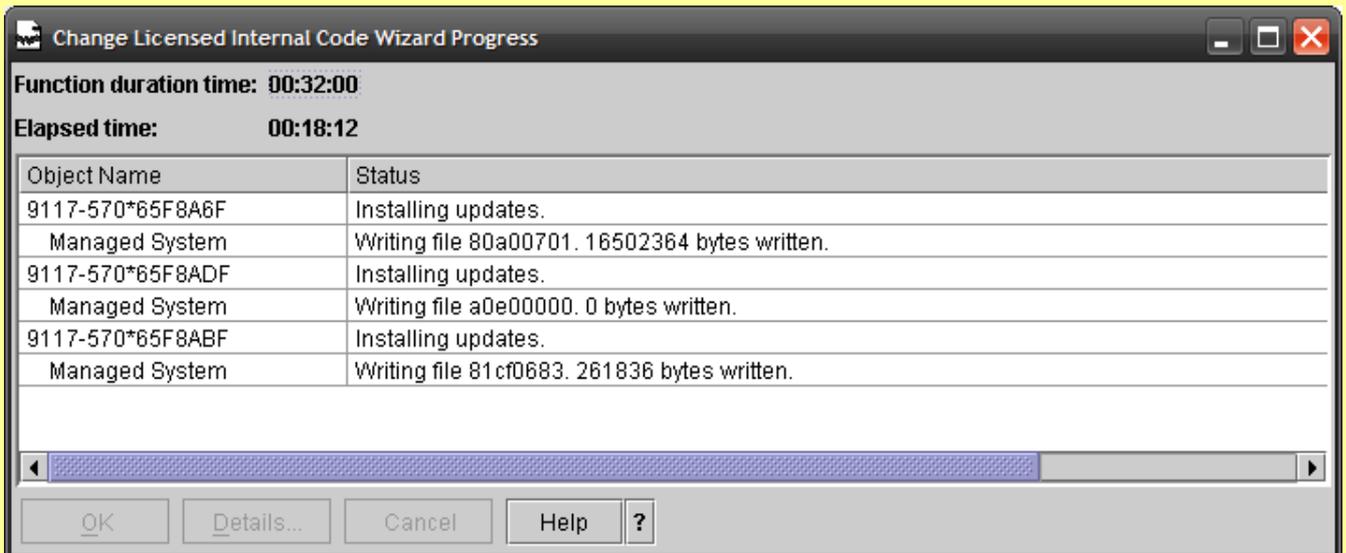


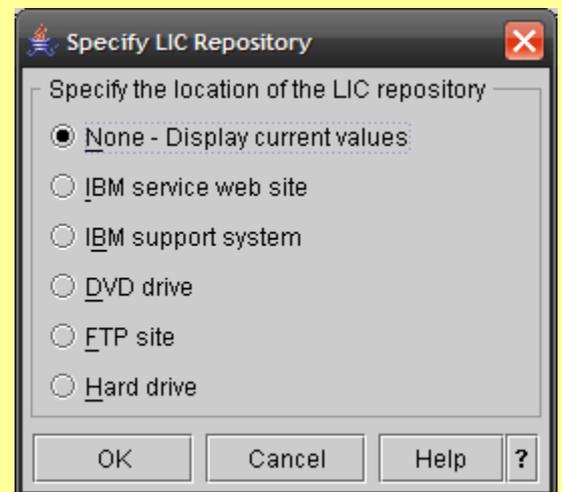
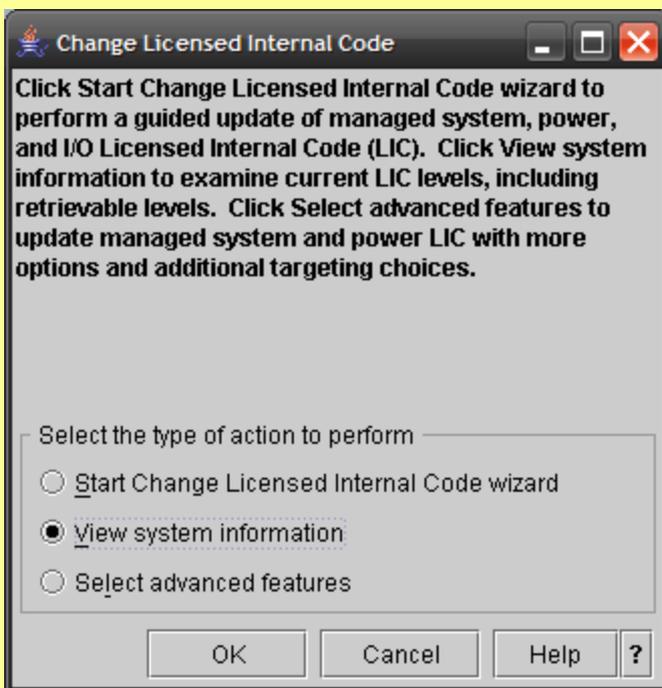
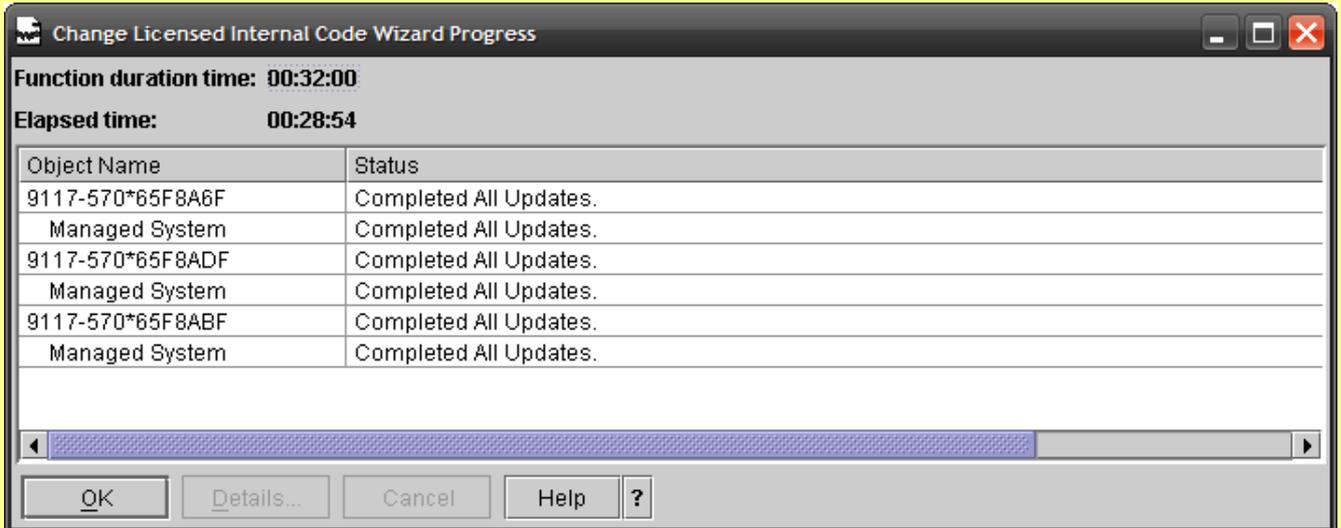
Here is the HMC check and confirmation that all upgrades are concurrent. Followed by this is to accept the license agreement.



If we keep navigating into the option, we can find system wise upgrade details, new fixes involved in each levels, intermittent to the final levels, concurrency state of the fix activation etc... We should not perform any configuration actions on the managed system as well as suggested to close the virtual terminal to lpars. If all environments are favorable, hmc update procedure automatically commits the current running firmware, a copy from temporary side to permanent side.







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

```

If the level of the microcode on the adapter is known, it is identified on the first line of the menu.

"M" is the most recent level of microcode found on the source.

It is later than the level of microcode currently installed on the adapter.

"L" identifies levels of microcode that are later than the level currently installed on the adapter.

If the current level is not known, all levels are considered later than the current level.

Multiple later images are listed in descending order.

"C" identifies the level of microcode currently installed on the adapter.

"P" identifies levels of microcode that are previous to the level currently installed on the adapter.

Multiple previous images are listed in the descending order.

```
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.

For most of the adapters to check the current microcode level for the following command can be used : *lscfg -vl <adapter device name>*.

Can also use the *lsmcode -A* command to check all the microcode versions.

To FLASH the EEPROM in the adapter using the single command: *diag -d <dev> -T download*

To back level the firmware : *diag -d <dev> -T "download -f -l previous"*

AIX Tuning.

All of us 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_rcvspace=262144
no -p -o tcp_sndspace=262144
no -p -o udp_rcvspace=655360
no -p -o udp_sndspace=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
```

```
### To restore the default ###
### ===== ###

no -p -D
vmo -p -D
ioo -p -D
```

```
### Only required on Servers running Databases ###
### And Databases using aio technology for IOs ###
##### AIO Configuration #####
#####=====#####
chdev -l aio0 -a -P autoconfig=available
chdev -l aio0 -a -P maxreqs=8192
chdev -l aio0 -a -P maxservers=100
chdev -l aio0 -a -P minservers=50
```

The IO parameter settings using ioo has to be done based on the IO pattern in the system. We may discuss this in one of the next versions.

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.

syslogd service.

syslogd service outputs are very useful for system/service/application troubleshooting.

Edit the file `/etc/syslog.conf` and add the following line at the end.

```
*.info /var/log/messages rotate size 1m files 4
```

Start/refresh the syslogd using `startsrc` and check the log file.

```
bash-3.00# startsrc -s syslogd
0513-029 The syslogd Subsystem is already active.
Multiple instances are not supported.
bash-3.00# refresh -s syslogd
0513-095 The request for subsystem refresh was completed successfully.
bash-3.00# cat /var/log/messages
cat: 0652-050 Cannot open /var/log/messages.
bash-3.00# touch /var/log/messages
bash-3.00# refresh -s syslogd
0513-095 The request for subsystem refresh was completed successfully.
bash-3.00# cat /var/log/messages
May 12 01:56:28 its syslog:info syslogd: restart
bash-3.00#
```

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 or beyond.

```
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 inl ine 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.

Atape Driver

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
----------------------	---	--------------------	------------------

[SAN Storage Drivers](#)

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.

<u>Driver support for DS6800 / DS8000 / SVC</u>			
FCP driver package	-	devices.fcp.disk.ibm.rte	V1.0.0.9
	-	devices.fcp.ibm2105.rte	V32.6.100.29
SDD driver package	-	devices.sdd.52.rte	V1.6.3.0
	-	devices.sdd.53.rte	V1.6.3.0

"devices.fcp.disk.ibm.rte" & "devices.fcp.ibm2105.rte" are host attachments package for AIX. The first package provides FCP device support for DS8000, DS6000, and SVC devices and the later one to ESS.

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)

Supported SDD features

The following SDD features are supported in the release 1.6.3.0 :

- 32- and 64-bit kernels
- Support for ESS, DS8000, DS6000 and virtualization products
- Preferred node path-selection algorithm for DS6000 and virtualization products.
- Changing the SDD path-selection algorithm dynamically.
 - Five path-selection algorithms are supported: –
 - Failover
 - Round robin
 - Round robin sequential
 - Load balancing
 - Load balancing sequential
- Dynamically adding paths to the SDD vpath devices
- Dynamically opening an invalid or close_dead path
- Dynamically removing or replacing PCI adapters or paths
- Fibre-channel dynamic device tracking v SDD server daemon support
- Support for HACMP v Support for secondary-system paging
 - Support for load-balancing and failover protection for AIX applications and LVM v SDD utility programs
- Support for SCSI-3 persistent reserve functions
- Support for AIX trace functions
 - Support more than 512 SAN Volume Controller devices from multiple SAN Volume Controller clusters on an AIX host
- Storage I/O priority feature in DS6000 and DS8000, only with AIX53 TL04 or later and with 64-bit kernel
- Virtual I/O Server with AIX 5.3 or later

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.

Installation of major files on your AIX host system :

The SDD installation package installs a number of major files on your AIX system.

defdpo	Define method of the SDD pseudo-parent data path optimizer (dpo).
cfgdpo	Configure method of the SDD pseudo-parent dpo.
define_vp	Define method of the SDD vpath devices.
addpaths	The command that dynamically adds more paths to SDD vpath devices while they are in <i>Available</i> state.
cfgvpath	Configure method of the SDD vpath devices.
chgvpath	Method to change vpath attributes.
cfallvpath	Fast-path configuration method to configure the SDD pseudo-parent dpo and all SDD vpath devices.
vpathdd	The SDD device driver.
hd2vp	The SDD script that converts an hdisk device volume group to an SDD vpath device volume group.
vp2hd	The SDD script that converts an SDD vpath device volume group to an hdisk device volume group.
datapath	The SDD driver console command tool.
lquerypr	The SDD driver persistent reserve command tool.
lsvpcfg	The SDD driver query configuration state command.
querysn	The SDD driver tool to query unique serial numbers of devices.
mkvg4vp	The command that creates an SDD volume group.
extendvg4vp	The command that extends the SDD vpath devices to an SDD volume group.
dpovgfix	The command that fixes an SDD volume group that has mixed vpath and hdisk physical volumes.
savevg4vp	The command that backs up all files belonging to a specified volume group with the SDD vpath devices.
restvg4vp	The command that restores all files belonging to a specified volume group with the SDD vpath devices.
sddsrv	The SDD server daemon for path reclamation and probe.
sample_sddsrv.conf	The sample SDD server configuration file.
lvmrecover	The SDD script that restores a system's SDD vpath devices and LVM configuration when a migration failure occurs.
sddfcmap	The SDD tool that collects information on ESS SCSI or disk storage systems fibre-channel devices through SCSI commands.
sddgetdata	The SDD data collection tool for problem analysis.

Is it OK to stop sddsrv in a live production machine ?

Answer is YES. Remember that hacmp failover or resource movement scripts do stop the sddsrv drivers for some time. The reason is that sddsrv as mentioned above is a daemon for path reclamation and probing. So it is not really affecting the storage IOs.

SDD migration procedure to follow in case of migrating the AIX OS.

- a. Run stopsrc -s sddsrv to stop the sddsrv daemon.
- b. Uninstall SDD.
- c. Upgrade to the latest version of the host attachment, if required.
The following are package names:
 - [ibm2105.rte](#) for 2105 devices
 - [devices.fcp.disk.ibm.rte](#) for 2145, 2107, and 1750 devices
- d. If rootvg is on a SAN boot disk, restart the system.
- e. Make sure no disk group is online except rootvg. Migrate the AIX OS level.
The system automatically restarts at the end of migration.
- f. Install SDD for the new AIX OS level.
- g. Configure SDD vpath devices by running the cfallvpath command.
- h. If required run the hd2vp command on all SDD volume groups.
- i. Resume all activities related to SDD devices.

Please do check the SDD User Guide for all other driver upgrade, or migration, hacmp migration etc..

Maximum number of LUNs :

For different AIX OS levels, SDD has set different limits on the maximum number of LUNs that can be configured. These limits exist because AIX has resource limitations on the total number of devices that a system can support. In a multipath configuration environment, AIX creates one hdisk device for each path to a physical disk. Increasing the number of paths that are configured to a physical disk increases the number of AIX system hdisk devices that are created and are consuming system resources. This might leave fewer resources for SDD vpath devices to be configured. On the other hand, more SDD vpath devices can be configured if the number of paths to each disk is reduced.

For AIX versions 4.3 and 5.1, AIX has a published limit of 10 000 devices per system. Based on this limitation, SDD limits the total maximum number of SDD vpath devices that can be configured to 600. This number is shared by all SDD-supported storage devices.

For AIX version 5.2 or later, the resource of the AIX operating system is increased. SDD has increased the SDD vpath device limit accordingly. Starting from SDD 1.6.0.7, SDD supports a combined maximum of 1200 supported storage devices on AIX version 5.2 or later.

OS level	SDD supported storage devices
AIX 4.3 *	600 LUNs (maximum 32 paths)
AIX 5.1	600 LUNs (maximum 32 paths)
AIX 5.2	1200 LUNs (maximum 32 paths; see Table 6 for recommended maximum number of paths.)
AIX 5.3	1200 LUNs (maximum 32 paths; see Table 6 for recommended maximum number of paths.)
AIX 6.1	1200 LUNs (maximum 32 paths; see Table 6 for recommended maximum number of paths.)

You can have a maximum of 32 paths per SDD vpath device regardless of the number of LUNs configured. However, configuring more paths than is needed for failover protection might consume too many system resources and degrade system performance. You should use the minimum number of paths necessary to achieve sufficient redundancy in the SAN environment. The recommended number of paths is 2 - 4.

Number of LUNs	Maximum paths per vpath
1- 600 vpath LUN	16
601 - 900 vpath LUN	8
901 - 1200 vpath LUN*	4

But it all again depends on the IO through put we require and the number of available HBA cards in the system.

Use the command-line interface to verify the configuration, enter `lsvpcfg`.

SAN boot support

With certain technology levels of the AIX operating systems, AIX supports fibre-channel boot capability for selected IBM System p. This allows you to select fibre-channel devices as the boot device. However, a multipathing boot device is not supported. If you plan to select a device as a boot device, do not configure that device with multipath configuration.

Please do remember to put off the SCSI reservation option on any of the boot disks. This may help in case of a system crash of the SAN boot systems due to the root disk access HBA failure. The reservation put by the failed HBA is hard to remove from the boot luns and this may prevent SMS routines from detecting the boot logical volume on this boot luns. AIX as such has no tool to remove a SCSI-2 reservation. The reservations put on a vpath is a SCSI-3 reservation and can be removed using `lquerypr` command.

chdev -l hdisk -a reservation_option=noreserve.

Ensure that the these boot luns are not shared with any other systems.

AIX data collection : (snap, perfpmr & sddgetdata)

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 <dir-full-path>*

The directory should be owned by user:group, "root:system".

```

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# █

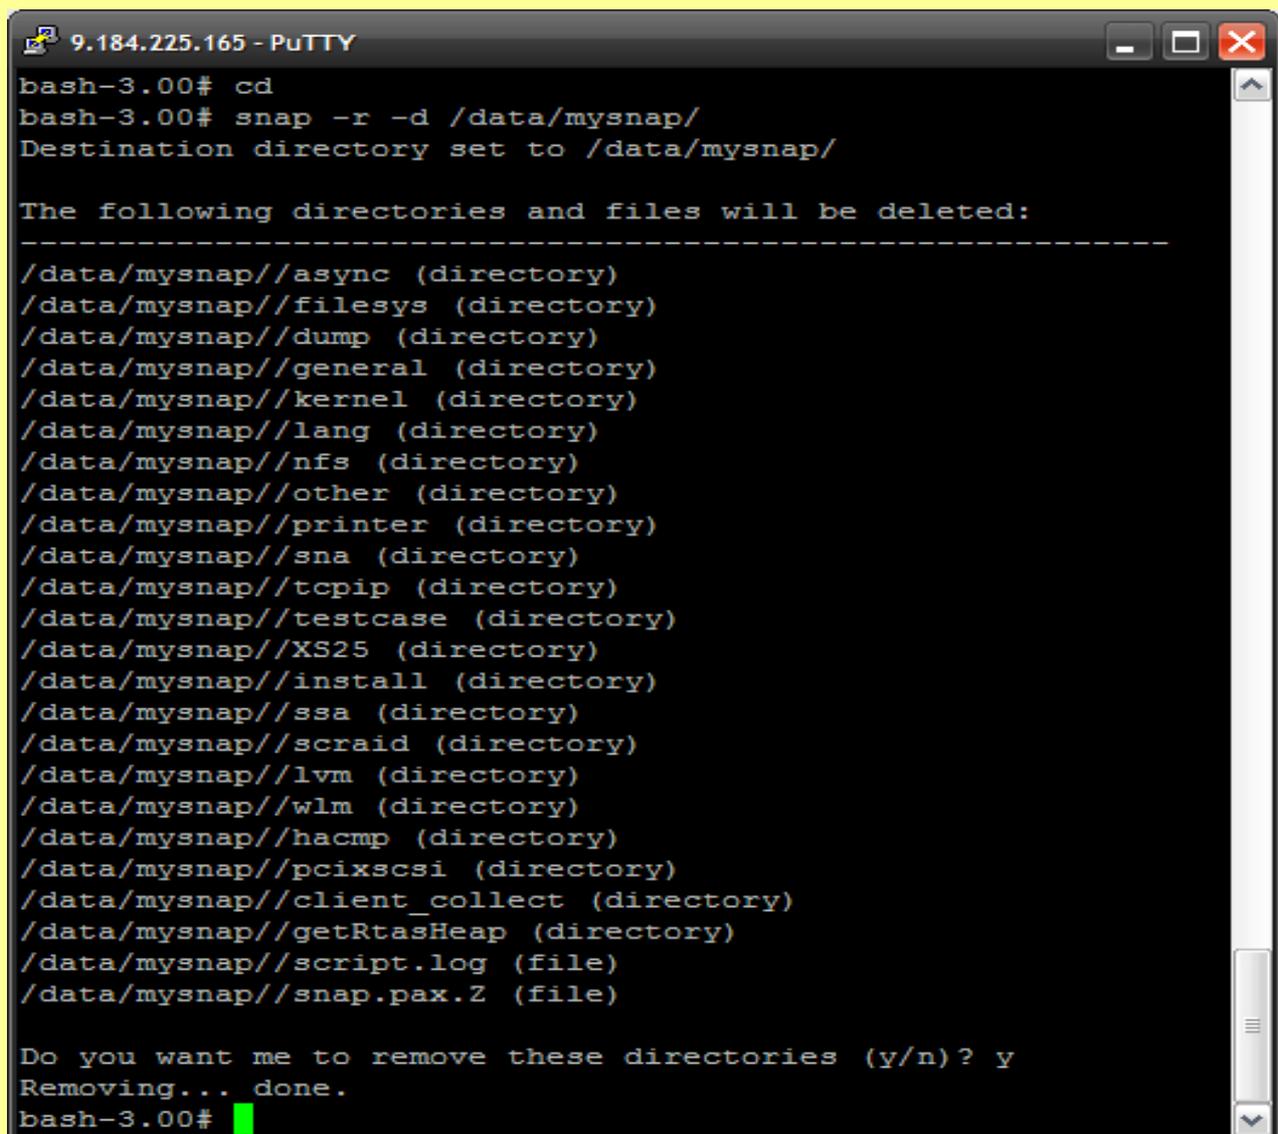
```

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//fileSYS (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 perfprm 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?

Using sddgetdata to collect information

SDD provides the **sddgetdata** script to collect information used for problem determination. For UNIX platforms, **sddgetdata** creates a tar file or a compressed tar file at the current directory with the current date and time as a part of the file name (for example, sdddata_hostname_yyyymmdd_hhmmss.tar or sdddata_hostname_yyyymmdd_hhmmss.tar.Z, where yyyymmdd_hhmmss is the timestamp of the file creation).

Follow the below given steps to collect the sddgetdata.

```
bash-3.00# pwd
/misc
bash-3.00# which sddgetdata
/usr/sbin/sddgetdata
bash-3.00# lslpp -w /usr/sbin/sddgetdata
File                               Fileset                               Type
-----
/usr/sbin/sddgetdata                devices.sdd.53.rte                   File
bash-3.00# /usr/sbin/sddgetdata
bash-3.00# pwd
/misc
bash-3.00# ls -l sdd*
-rw-r--r--  1 root      system      11438080 May 24 04:37 sdddata_drfinr1_20080524_043727.tar
```

It collects the following for files which can be used for further analysis.

```
bash-3.00# ls -l
total 22384
-rw-r--r--  1 root      system      29012 May 24 04:32 datapath.20080524_043212
-rw-r--r--  1 root      system     2235690 May 24 04:32 errpt.20080524_043212
-rw-r--r--  1 root      system        66 May 23 20:33 lvm.cfg
-rw-r--r--  1 root      system      516 May 23 20:34 post_i.out
-rw-r--r--  1 root      system      358 Oct 26 2006 post_u.out
-rw-r--r--  1 root      system     1123 May 23 20:33 pre_rm.out
-rw-r--r--  1 root      system      890 Oct 26 2006 pre_u.out
-rw-r--r--  1 root      system     457917 May 24 04:26 sdd.log
-rw-r--r--  1 root      system     3932407 May 17 16:13 sdd_bak.log
-rw-r--r--  1 root      system      1129 May 24 04:32 sddinfo.20080524_043212
-rw-r--r--  1 root      system     570494 May 24 04:32 sddsrv.log
-rw-r--r--  1 root      system     4194351 May 23 15:23 sddsrv_bak.log
```

datapath.20080524_043212 : datapath query adapter and datapath query device output

errpt.20080524_043212 : copy of errpt -a output

lvm.cfg : copy of /etc/lvm.cfg file.

post_i.out

post_u.out

pre_rm.out

pre_u.out

sdd.log : copy of /var/adm/ras/sdd.log

sdd_bak.log : copy of /var/adm/ras/sdd_bak.log

sddinfo.20080524_043212 : sdd and fcp version information and the dpo device attributes.

sddsrv.log : copy of /var/adm/ras/sddsrv.log

sddsrv_bak.log : copy of /var/adm/ras/sddsrv_bak.log

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.

NIM : AIX Installation & Configuration

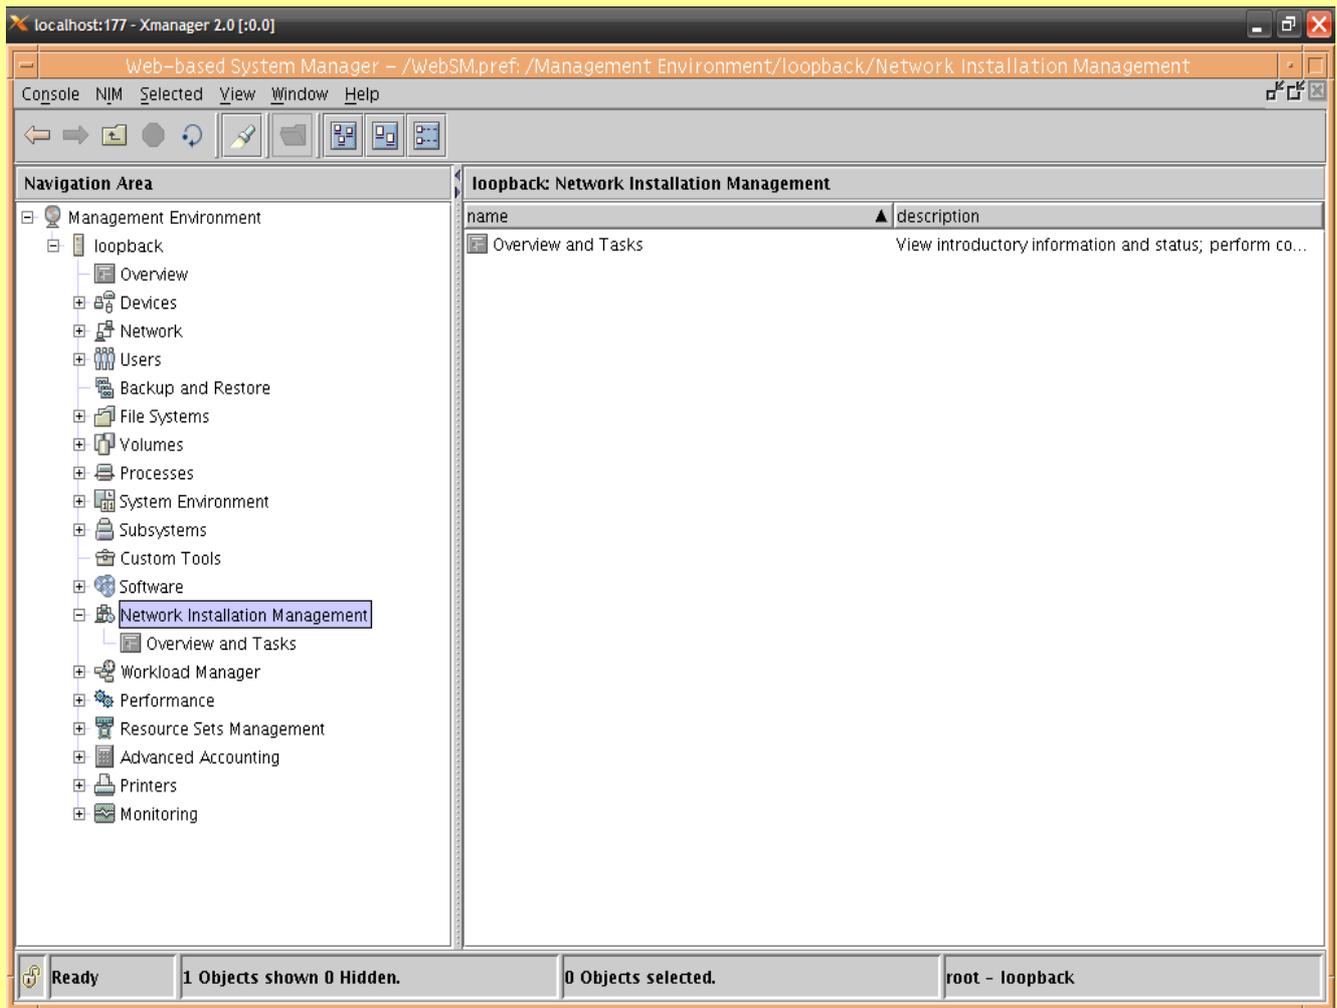
NIM enables system administrator to centrally manage the installation and configuration of AIX and optional software on machines within a network environment. Setting up NIM involves various tasks including the following:

- Installing NIM file sets,
- Configuring basic resources,
- Creating machine and network definitions
- Creating resources that are used to install the nodes.

The specific tasks that you need to perform depend on which features of NIM that you plan to use. For more information about NIM, see the *IBM AIX Installation Guide and Reference* depending on the version of AIX that you are using (AIX 5L 5.2 or 5.3).

This below discussions are very basic setup information related to NIM. Graphical method is preferred as it is easy to work with for beginners. NIM configuration, is one of the unique scenarios, where I personally prefer to login to X windows. Some how the comfort lies in using the graphics for NIM deployments.

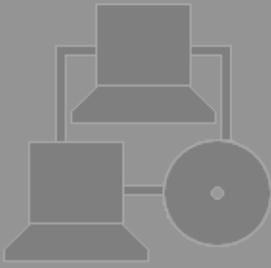
Starting by logging into the X windows subsystem, you may have to run /etc/rc.dt to start the X windows in AIX, open a terminal and run the command “wsm”. I used X manager package from laptop for X login. Most of the screens given below are self informative.



Navigate to the “Network Installation Management” option and select the Overview & Tasks sub menu.

Network Installation Management: Overview and Tasks

Network Installation Management (NIM)



Network Installation Management (NIM) supports the installation of AIX across a network. One machine in a NIM environment must be designated as the master to control the installation of NIM clients. All clients must be defined in the master's database before software can be installed on them.

[More Information](#)

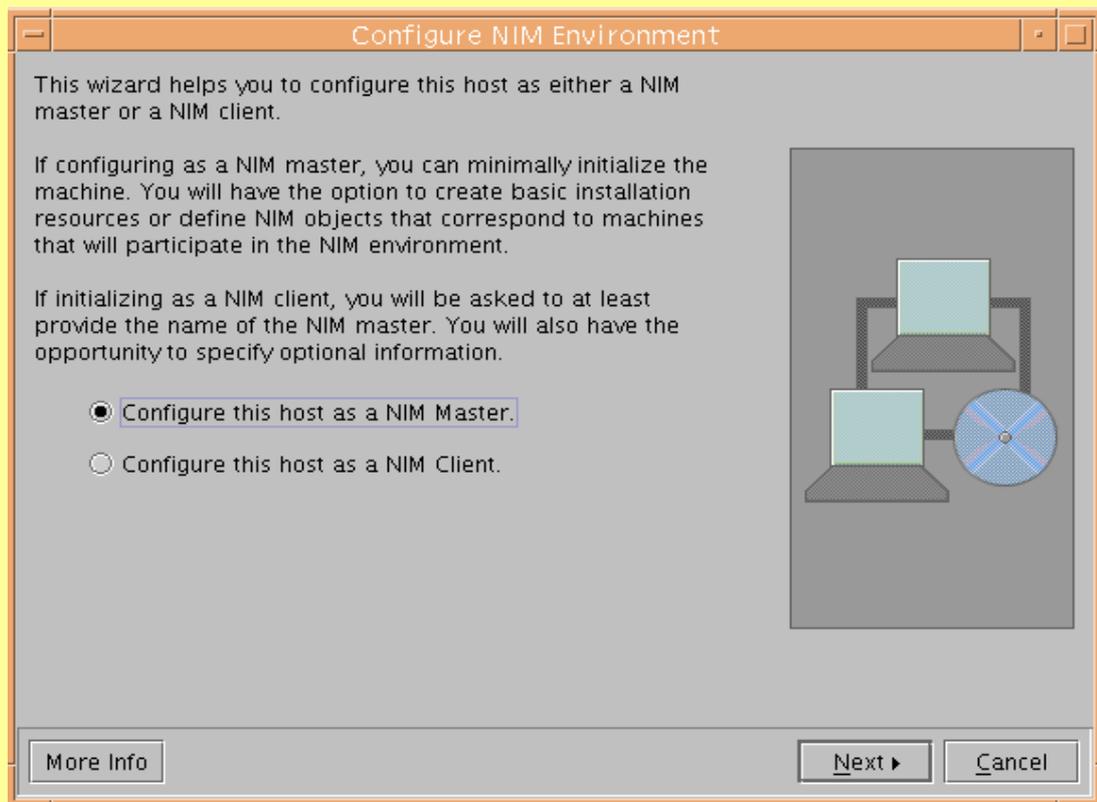
[Configure the Network Installation Management environment](#)

[Configure Client Communication Services](#)

TASKS

STATUS

Network Installation Management environment:	Unconfigured
Number of Machines defined:	0
Number of Machines Running:	0
Number of Machines being Processed:	0



Configuring the machine as NIM master may fail as the required bos package is not present.

```
bash-3.00# lslpp -l | grep nim
X11.Dt.helpmin          5.3.0.0  COMMITTED  AIX CDE Minimum Help Files
X11.msg.en_US.Dt.helpmin 5.3.0.0  COMMITTED  AIX CDE Minimum Help Files -
bos.sysmgt.nim.client    5.3.0.63 APPLIED    Network Install Manager -
X11.Dt.helpmin          5.3.0.0  COMMITTED  AIX CDE Minimum Help Files
bos.sysmgt.nim.client    5.3.0.63 APPLIED    Network Install Manager -
```

```
bash-3.00# oslevel -s
5300-06-05-0806
```

System need to have nim.master package to proceed further.

The package can be installed from the AIX BOS Package CDs as follows.

The BOS Package CD used is AIX Version 5.3 TL04 where as the running version is 5.3 TL06+SP5.

Hence followed by the base package installation sufficient patching must be done to bring the system back to the running level.

```

                                Install Software

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

                                [Entry Fields]
* INPUT device / directory for software      .
* SOFTWARE to install                        [+ 5.3.0.40 Network Install Manager - Master Tools > +
PREVIEW only? (install 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 +
OVERWRITE same or newer versions?          no +
VERIFY install and check file sizes?        no +
Include corresponding LANGUAGE filesets?    yes +
DETAILED output?                            no +
Process multiple volumes?                   yes +
ACCEPT new license agreements?              [no] +
PREVIEW new LICENSE agreements?             no +

```

First we tried to install from a dump of usr/sys/inst.images, but it failed as the nim.master package is in the location install/ppc/.

```
installp: APPLYING software for: bos.sysmgt.nim.master 5.3.0.40
```

```
cannot open /swdump/AIX_53_TL4_BASE/installp/ppc/bos.sysmgt: No such file or directory
Please mount volume 1 on /swdump/AIX_53_TL4_BASE/installp/ppc/bos.sysmgt
...and press Enter to continue installp: An error occurred while running the restore command.
Use local problem reporting procedures.
```

```
installp: CANCELLED software for: bos.sysmgt.nim.master 5.3.0.40
```

We used CD instead...

```

Command: running          stdout: yes          stderr: no

Before command completion, additional instructions may appear below.

[MORE...22]
<< End of Success Section >>

+-----+
+          BUILDDATE Verification ...          +
+-----+
Verifying build dates...done
FILESET STATISTICS
-----
  1 Selected to be installed, of which:
    1 Passed pre-installation verification
  ----
  1 Total to be installed

+-----+
+          Installing Software...          +
+-----+

installp: APPLYING software for:
          bos.sysmgt.nim.master 5.3.0.10

. . . . . << Copyright notice for bos.sysmgt >> . . . . .
Licensed Materials - Property of IBM

5765G0300
(C) Copyright International Business Machines Corp. 1993, 2004.

All rights reserved.
US Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corp.
. . . . . << End of copyright notice for bos.sysmgt >>. . . . .

```

```

All rights reserved.
US Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corp.
. . . . . << End of copyright notice for bos.sysmgt >>. . . . .

Finished processing all filesets. (Total time: 27 secs).

+-----+
+          Summaries:          +
+-----+

Installation Summary
-----
Name                          Level          Part           Event          Result
-----
bos.sysmgt.nim.master          5.3.0.10      USR            APPLY          SUCCESS

File /etc/inittab has been modified.

One or more of the files listed in /etc/check_config.files have changed.
See /var/adm/ras/config.diff for details.

```

```

# oslevel -s
5300-01-00-0000
# oslevel -sq
Known Service Packs
-----
5300-06-05-0806
5300-06-04-0748
5300-06-03-0732
5300-06-02-0727
5300-06-01-0722
5300-05-CSP-0000
5300-05-06-0000
5300-05-05-0000
5300-05-04-0000
5300-05-03-0000
5300-05-02-0000
5300-05-01-0000
5300-04-CSP-0000
5300-04-03-0000
5300-04-02-0000
5300-04-01-0000
5300-03-CSP-0000
5300-01-00-0000
# oslevel -sl 5300-06-05-0806
Fileset                                Actual Level      Service Pack Level
-----
bos.sysmgt.nim.master                  5.3.0.10          5.3.0.61

```

Trying to bring back system to its running level by re-patching. Used smitty update_all with TL06 patch

```

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?                yes
PREVIEW new LICENSE agreements?              no

```

Then used smitty update_all with TL06+SP05

```

# oslevel -s
5300-06-03-0732
# oslevel -sq
Known Service Packs
-----
5300-06-05-0806
5300-06-04-0748
5300-06-03-0732
5300-06-02-0727
5300-06-01-0722
5300-06-00-0000
5300-05-CSP-0000
5300-05-06-0000
5300-05-05-0000
5300-05-04-0000
5300-05-03-0000
5300-05-02-0000
5300-05-01-0000
5300-04-CSP-0000
5300-04-03-0000
5300-04-02-0000
5300-04-01-0000
5300-03-CSP-0000
# oslevel -sl 5300-06-05-0806
Fileset                                Actual Level      Service Pack Level
-----
bos.sysmgt.nim.master                  5.3.0.60          5.3.0.61

```

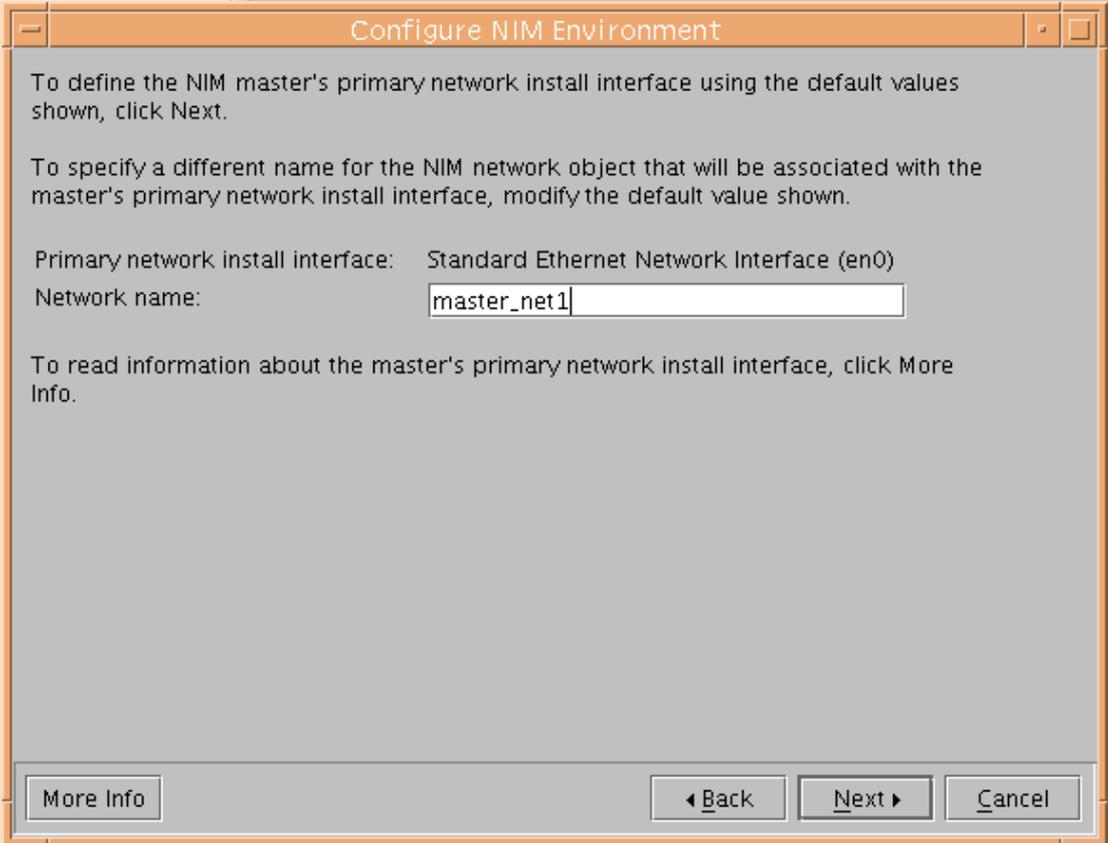
```

# oslevel -s
5300-06-05-0806
# oslevel -sl 5300-06-05-0806
# lslpp -l | grep nim
X11.Dt.helpmin                        5.3.0.0   COMMITTED   AIX CDE Minimum Help Files
X11.msg.en_US.Dt.helpmin              5.3.0.0   COMMITTED   AIX CDE Minimum Help Files -
bos.sysmgt.nim.client                 5.3.0.63   APPLIED     Network Install Manager -
bos.sysmgt.nim.master                 5.3.0.61   COMMITTED   Network Install Manager -
X11.Dt.helpmin                        5.3.0.0   COMMITTED   AIX CDE Minimum Help Files
bos.sysmgt.nim.client                 5.3.0.63   APPLIED     Network Install Manager -
#

```

After ensuring that nim.master package is installed and system back to its originally patched level, continue the NIM configuration from wsm.

We are more interested to configure the system with all custom configuration as against the default. We had already created a filesystem and mounted it under /nim. Intention is to keep all the nim resources under the /nim filesystem.



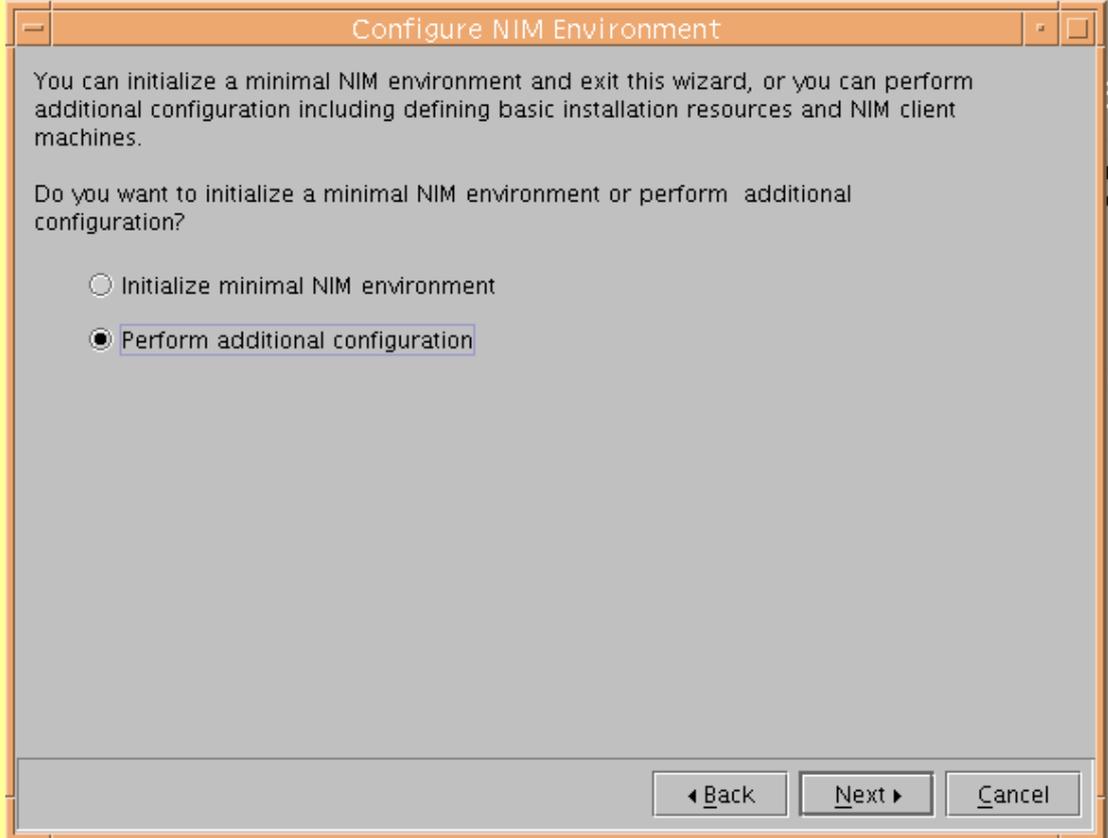
Configure NIM Environment

To define the NIM master's primary network install interface using the default values shown, click Next.

To specify a different name for the NIM network object that will be associated with the master's primary network install interface, modify the default value shown.

Primary network install interface: Standard Ethernet Network Interface (en0)
Network name:

To read information about the master's primary network install interface, click More Info.



Configure NIM Environment

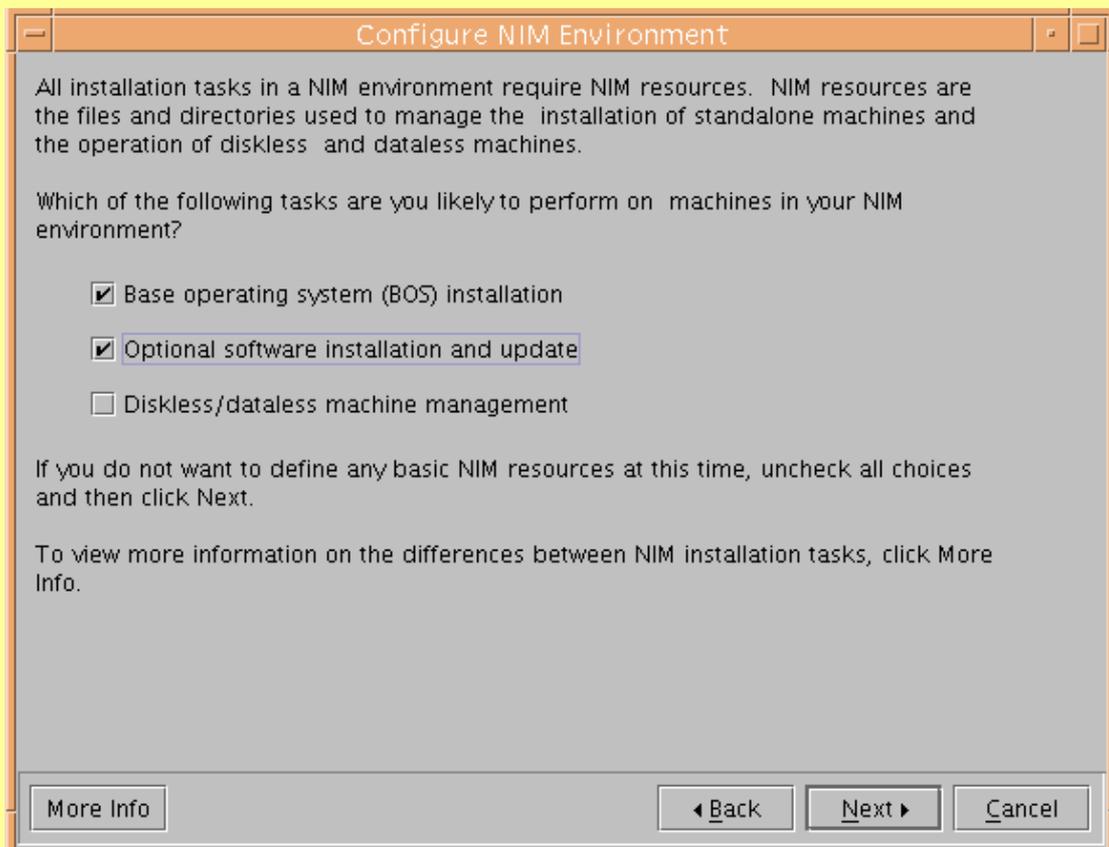
You can initialize a minimal NIM environment and exit this wizard, or you can perform additional configuration including defining basic installation resources and NIM client machines.

Do you want to initialize a minimal NIM environment or perform additional configuration?

Initialize minimal NIM environment
 Perform additional configuration

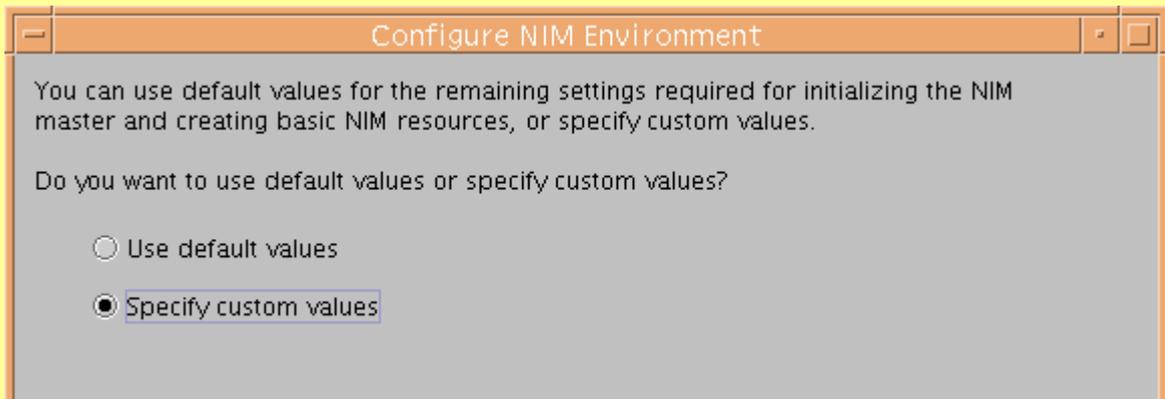
```
# df -g
Filesystem      GB blocks      Free %Used      Iused %Iused Mounted on
/dev/hd4         1.00          0.96   5%         2881    2% /
/dev/hd2         4.00          2.12  47%        40841   8% /usr
/dev/hd9var      4.00          3.98   1%          470    1% /var
/dev/hd3         4.00          3.99   1%          103    1% /tmp
/dev/hd1         1.00          1.00   1%           5     1% /home
/proc            -             -       -           -     - /proc
/dev/hd10opt     1.00          0.90  10%         3088    2% /opt
/dev/swdump     15.00         3.18  79%         5307    1% /swdump
/dev/nim         6.00          3.77  38%           5     1% /nim
```

Generally a 4GB filesystem is well enough.

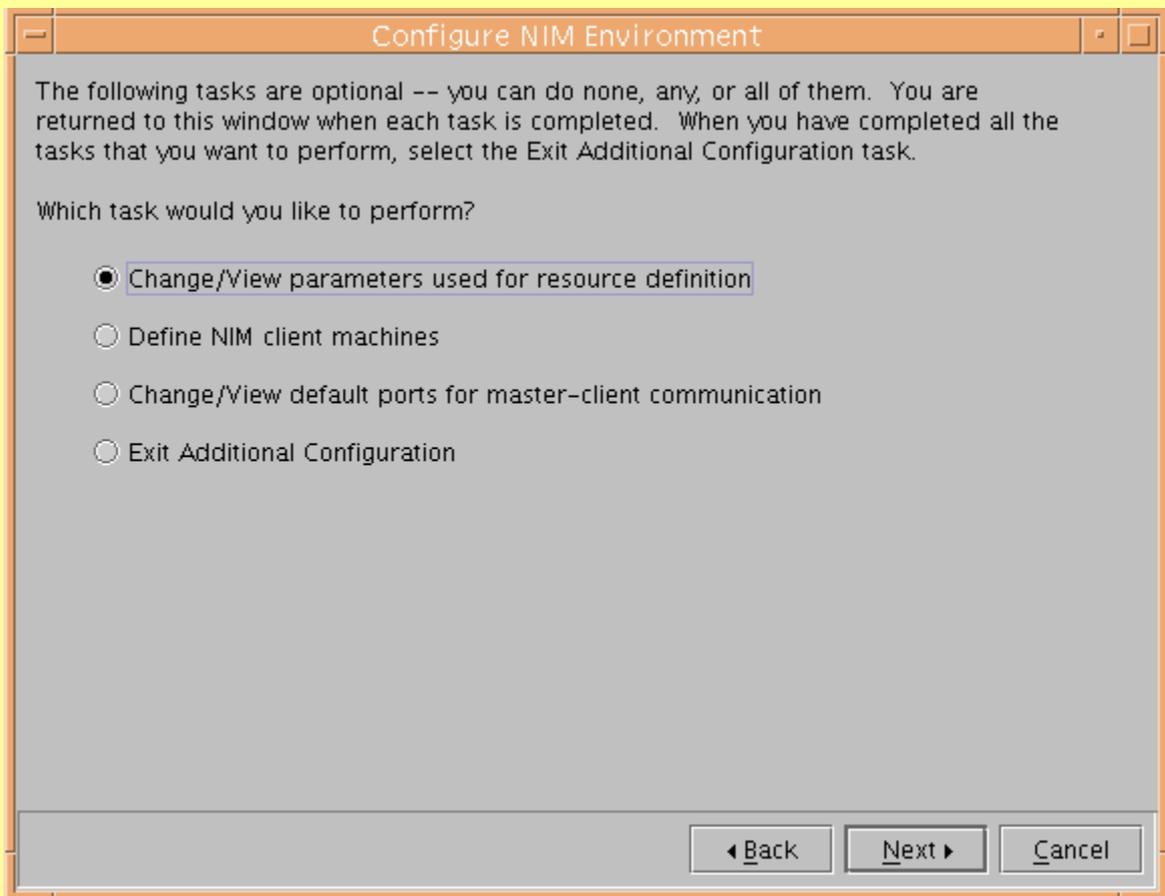


We also copied all the AIX packages to /swdump/AIX_53_TL4_BASE directory.

Our idea is to finally deploy AIX 5.3 with TL06 + SP5. But we had to start with the TL04 Base CD as we are not having a TL06 Base CD.



We opted for custom values and changed the parameters as shown below.



Configure NIM Environment

Verify or modify the lpp_source characteristics shown.

NIM name:

Location:

Comments:

Create a file system for the lpp_source images

File system size (megabytes):

Volume group for new file system:

To view further information on lpp_source creation click [More Info](#).

Configure NIM Environment

The SPOT resource contains key software required to boot and install machines over the network. It can be created in its own directory or file system, or it can be combined with the software already installed in the /usr file system of the server.

Do you want to create a SPOT in a new directory or combine its software with the system's installed software?

Create SPOT in new directory or file system

Combine SPOT software with system's installed software

Configure NIM Environment

Verify or modify the SPOT characteristics shown.

NIM name:

Location:

Comments:

Create a file system for the SPOT resource

File system size (megabytes):

Volume group for new file system:

Configure NIM Environment

Bundles are collections of packages or individual filesets that suit a particular purpose. By default, this wizard defines NIM bundle resources that are equivalent to the system-defined software bundles in the `/usr/sys/inst.data/sys_bundles` directory.

Do you want to create the system-defined bundle resources?

Create system-defined bundle resources

Do not create system-defined bundle resources

Configure NIM Environment

The following tasks are optional -- you can do none, any, or all of them. You are returned to this window when each task is completed. When you have completed all the tasks that you want to perform, select the Exit Additional Configuration task.

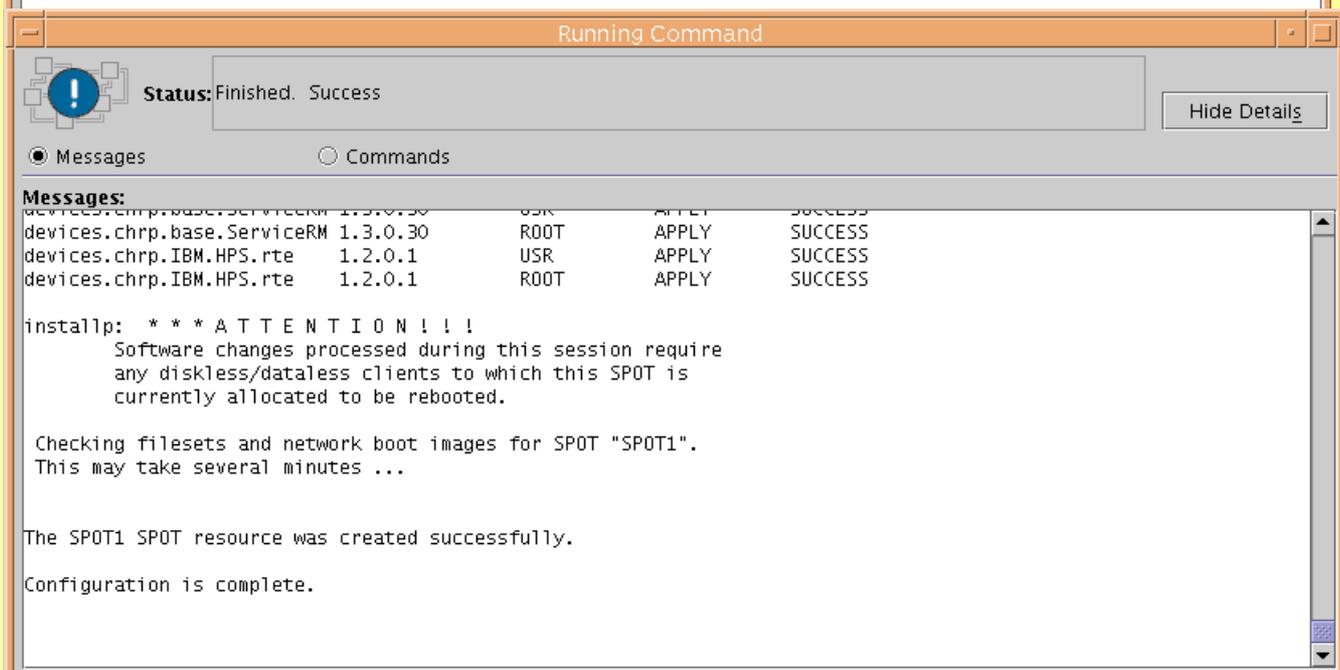
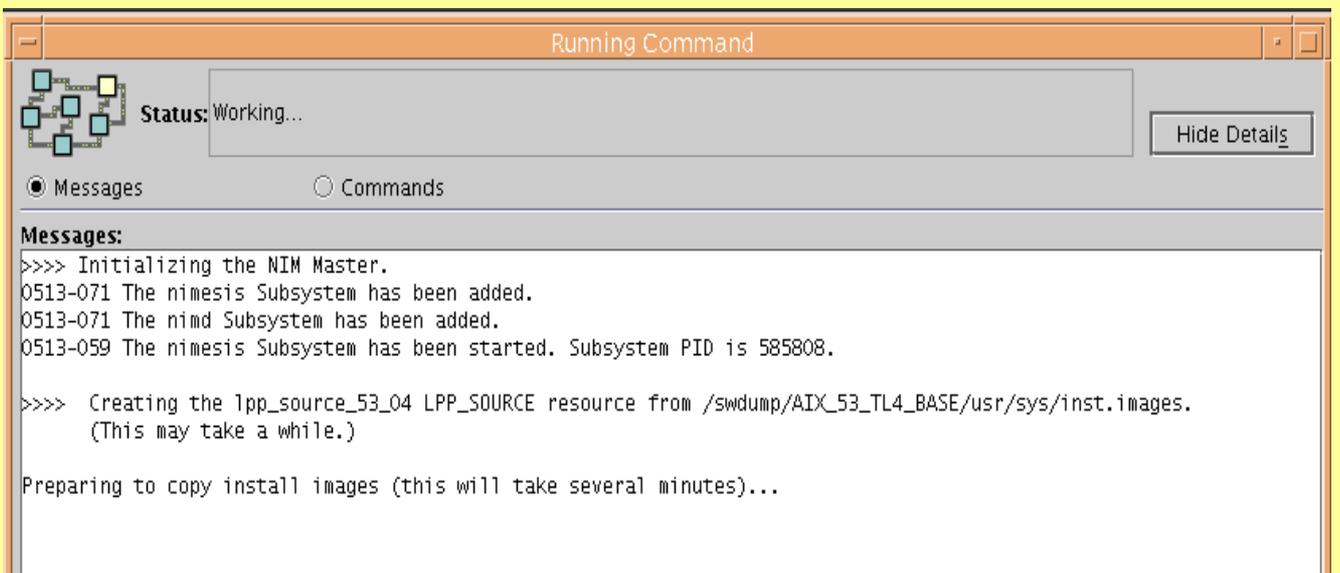
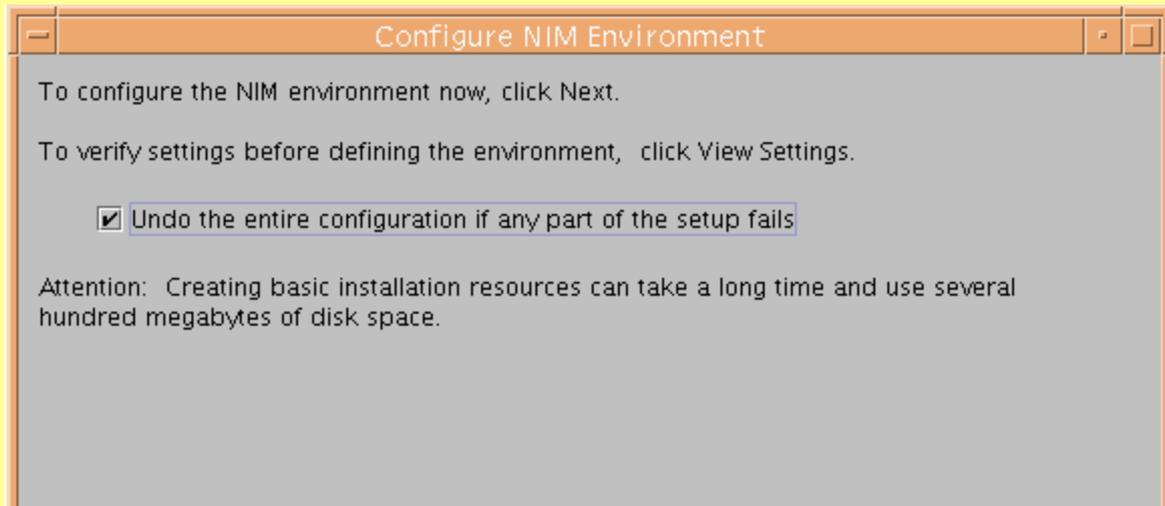
Which task would you like to perform?

Change/View parameters used for resource definition

Define NIM client machines

Change/View default ports for master-client communication

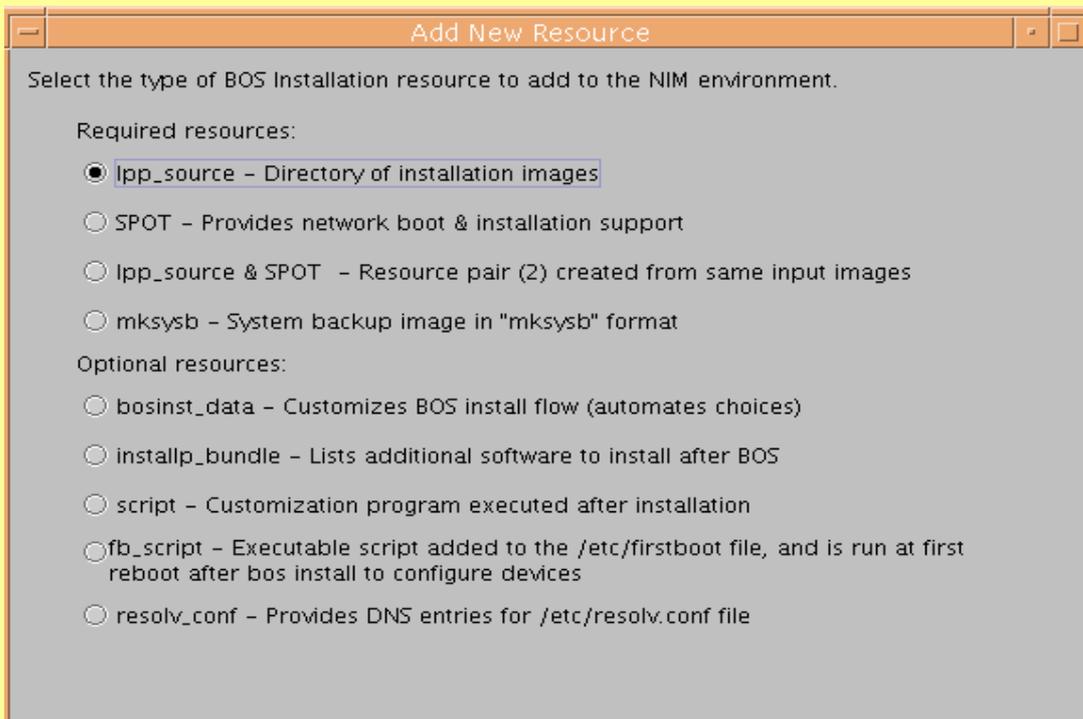
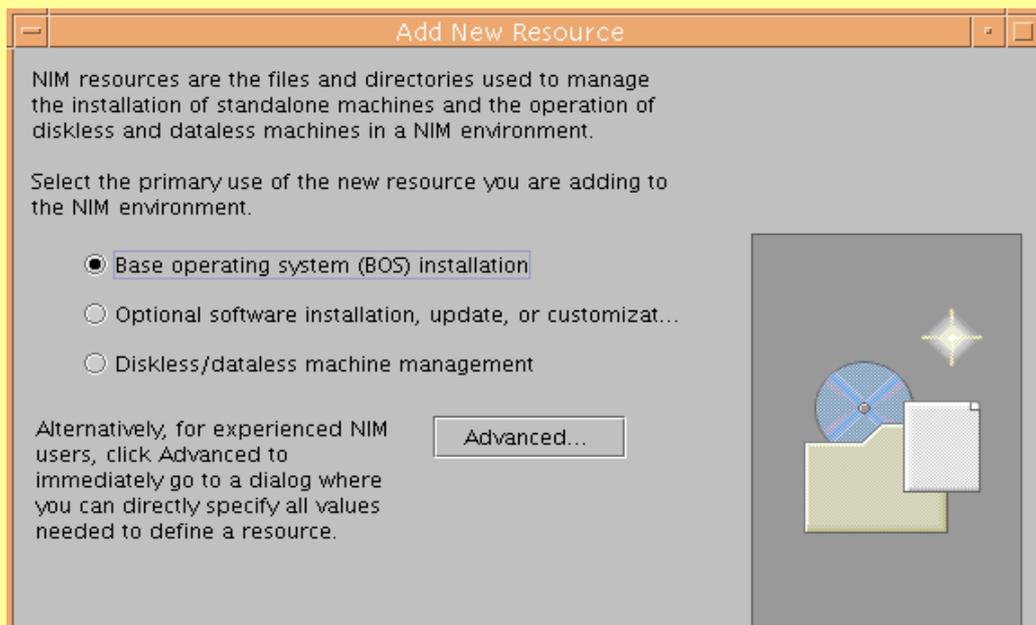
Exit Additional Configuration



Network Installation Management: Resources				
Name	Type	State	Additional Information	
 lpp_source_53_04	lpp_source	ready for use	LPP_53_TL04	
 SPOT1	spot	ready for use	SPOT_53_TL04	

Now we have a lpp resource and a SPOT resource ready. SPOT is fundamentally used as a tftp image to bootup and lpp is the installation program bundle.

Next is to create two more lpp resources, one with TL06 and another one with TL06_SP5 patches. Use the “Add New Resource” option from the Resources option in the NIM.



Add New Resource

An lpp_source resource can be created from software images on a CD or tape device, from an existing lpp_source, or from images currently residing in a directory.

Select the preferred method of creation, and then select or specify the software image source.

Create resource from device or existing resource

16 Bit LVD SCSI DVD-RAM Drive (cd0)

Create resource from directory

/swdump/AIX_53_TL6/

Select the architecture type of the source images to copy over when you create the LPP source. If no architecture is specified, an attempt is made to determine the architecture.

Architecture Type: power

Add New Resource

You can use default values for the remaining settings required to define the resource, or specify custom values.

Do you want to use default values or specify custom values?

Use default values

Specify custom values

Add New Resource

Verify or modify the lpp_source characteristics shown.

NIM name: lpp_source_53_06

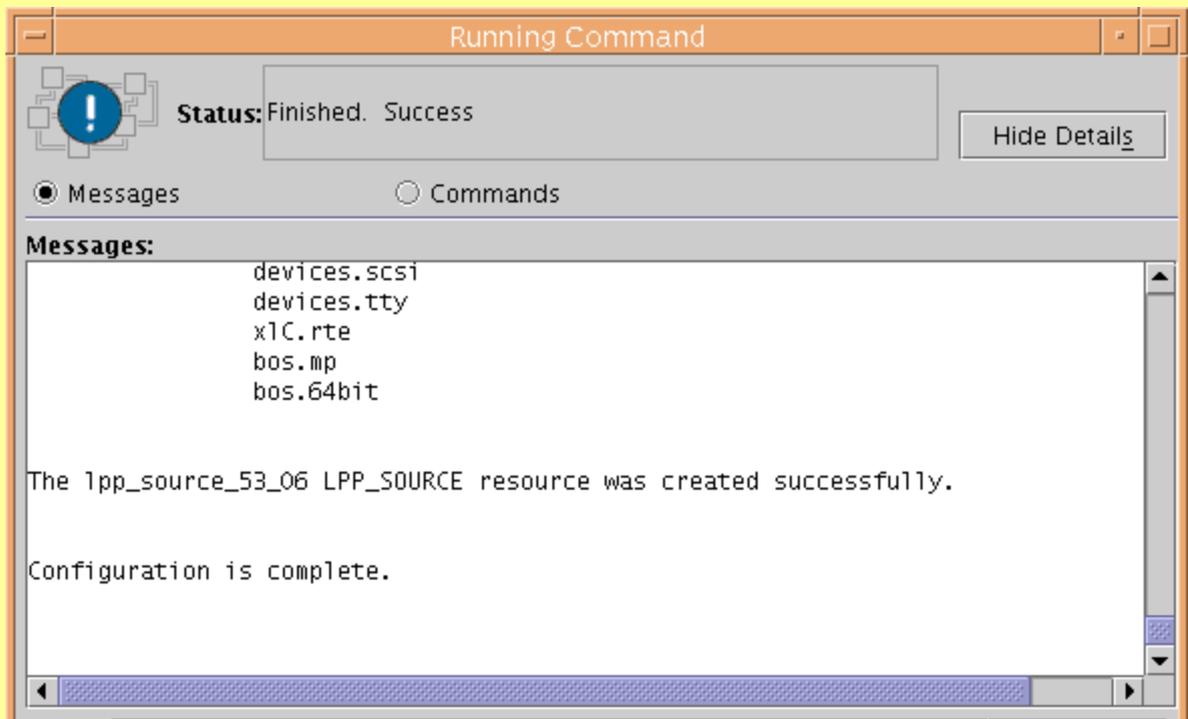
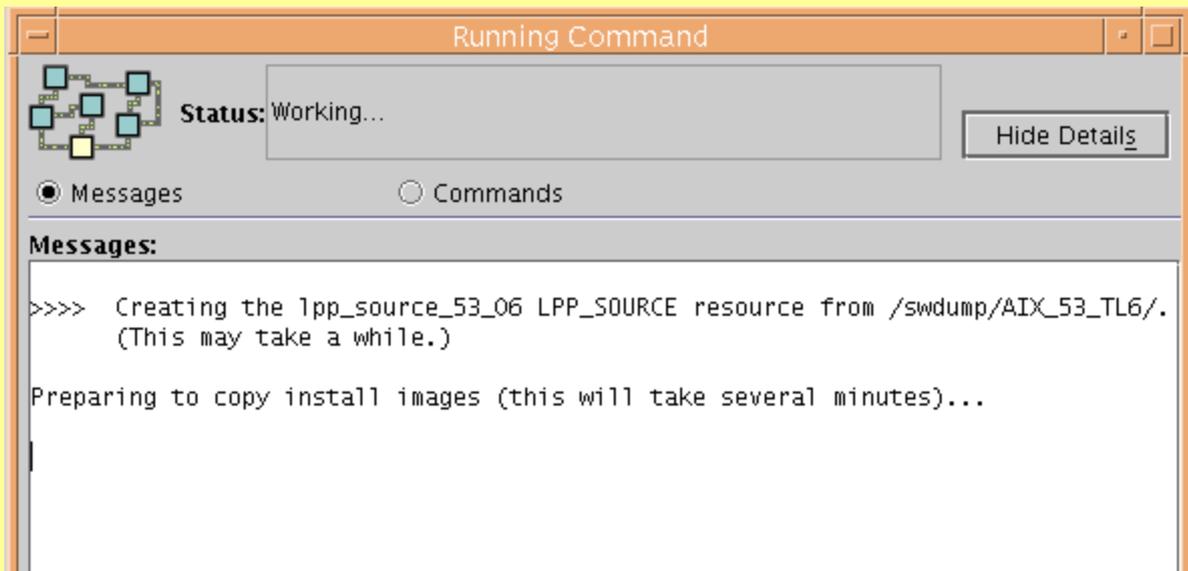
Path: /nim/lpp_source_53_06

Comments: LPP_53_TL06-Only Patch

Create a file system for the lpp_source images

File system size (megabytes): 650

Volume group for new file system: rootvg



Next is to create TL06_SP5 lpp.

Verify or modify the lpp_source characteristics shown.

NIM name:

Path:

Comments:

Create a file system for the lpp_source images

File system size (megabytes):

Volume group for new file system:

An lpp_source resource can be created from software images on a CD or tape device, from an existing lpp_source, or from images currently residing in a directory.

Select the preferred method of creation, and then select or specify the software image source.

Create resource from device or existing resource

Create resource from directory

16 Bit LVD SCSI DVD-RAM Drive (cd0)

Select the architecture type of the source images to copy over when you create the LPP source. If no architecture is specified, an attempt is made to determine the architecture.

Architecture Type:

Status: Finished. Success

Messages Commands

Messages:

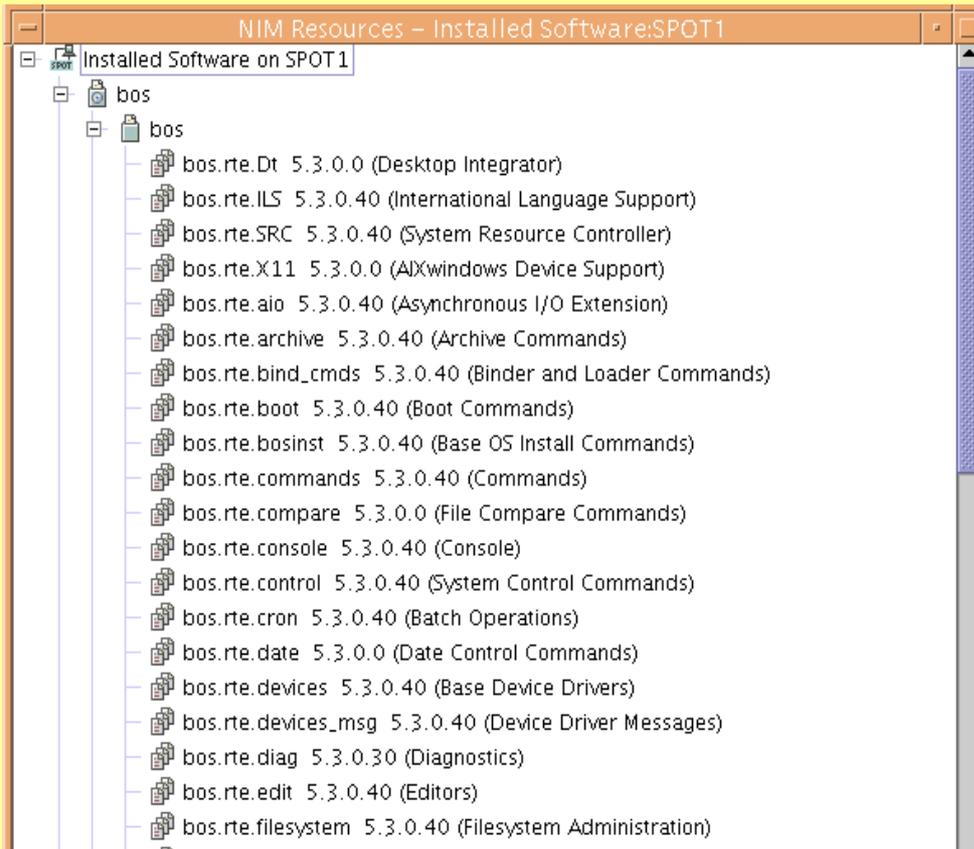
```

devices.graphics
devices.scsi
devices.tty
x1C.rte
bos.mp
devices.common
bos.64bit

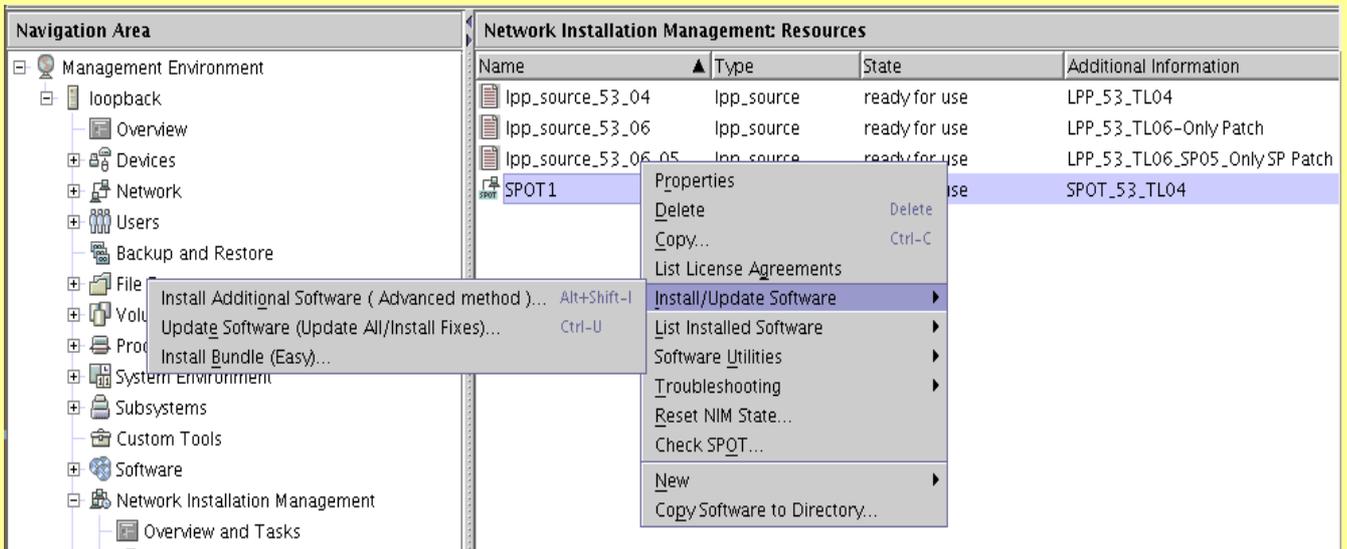
The lpp_source_53_06_05 LPP_SOURCE resource was created successfully.

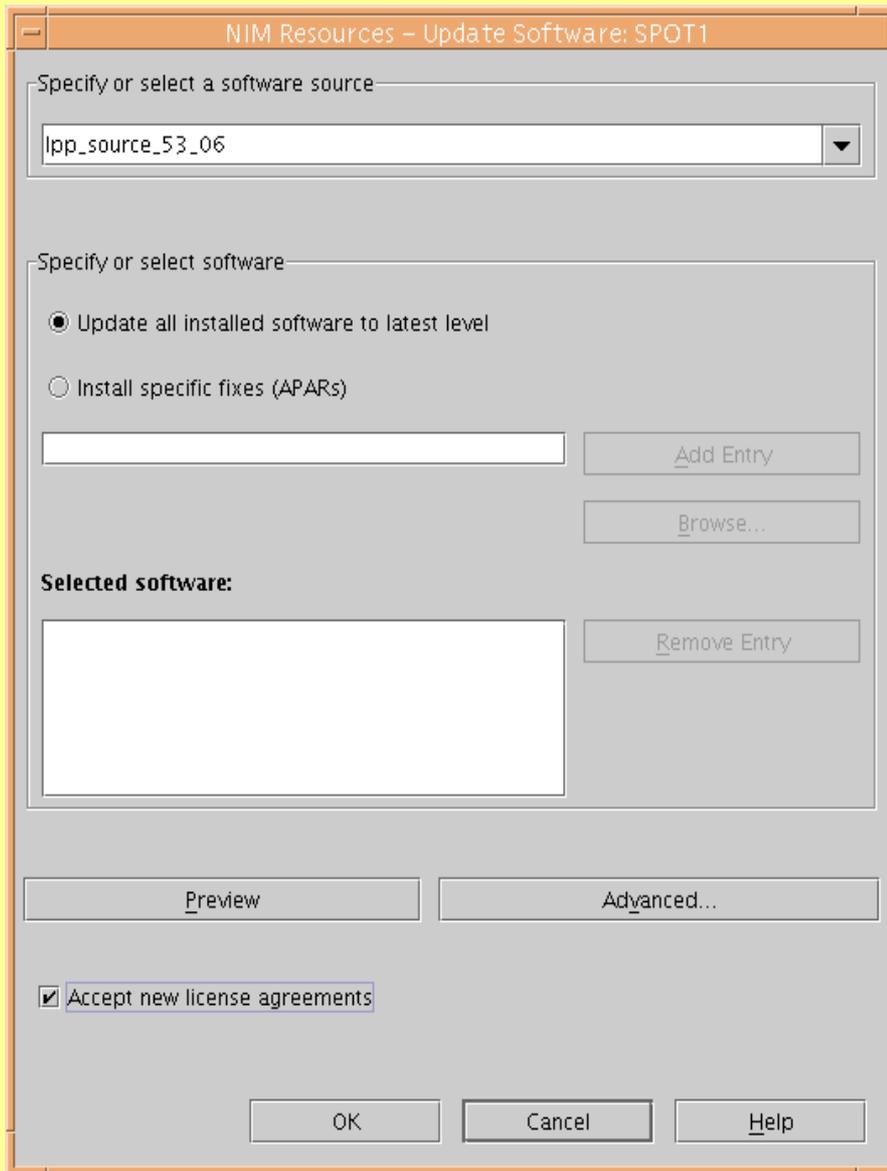
Configuration is complete.
```

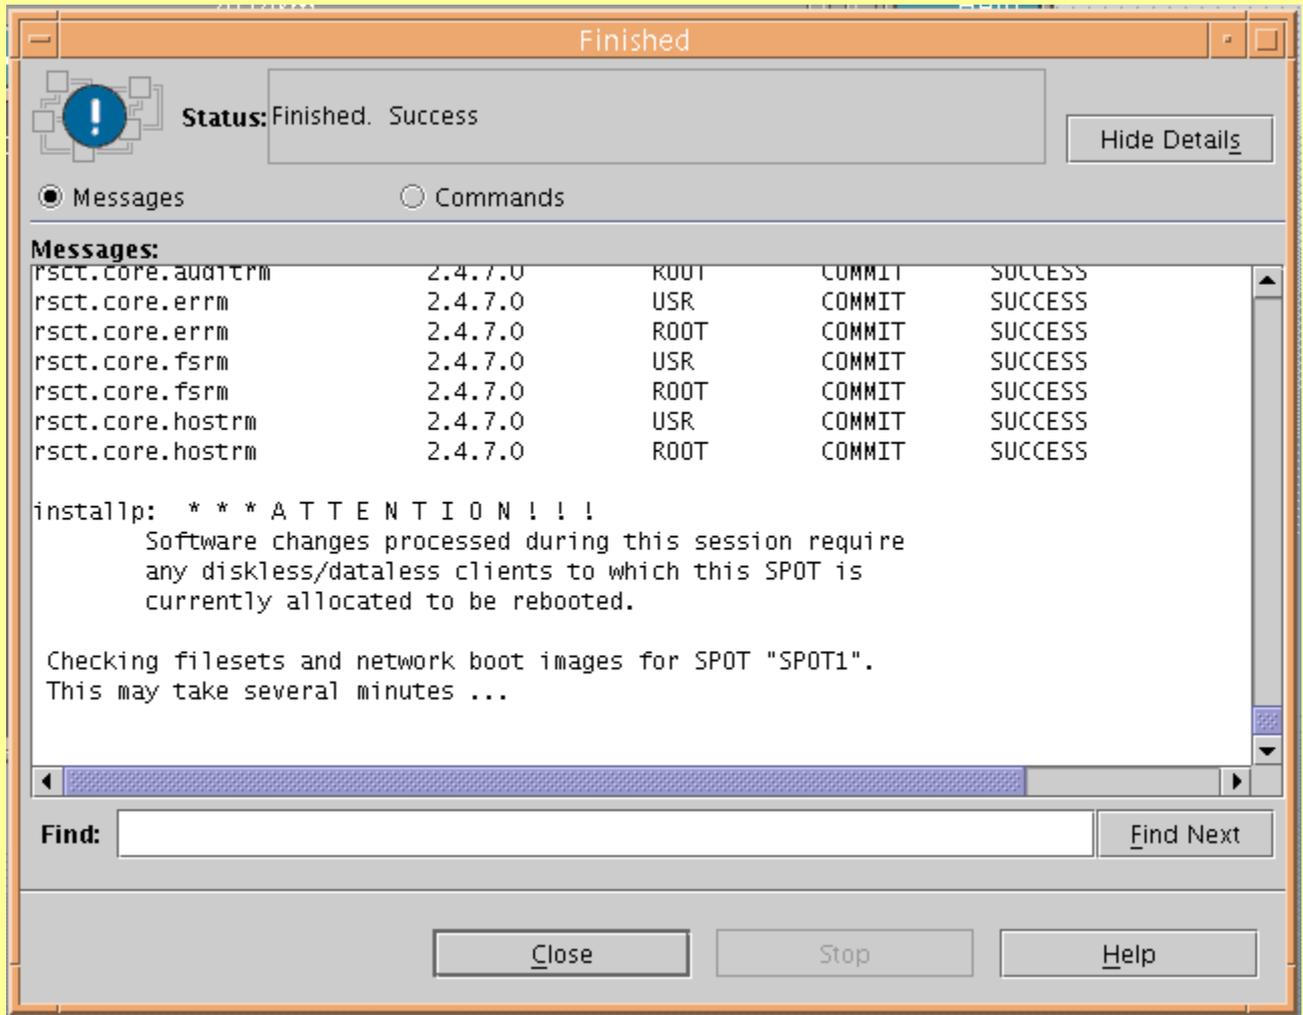
Checking the software package information in each of these resources.

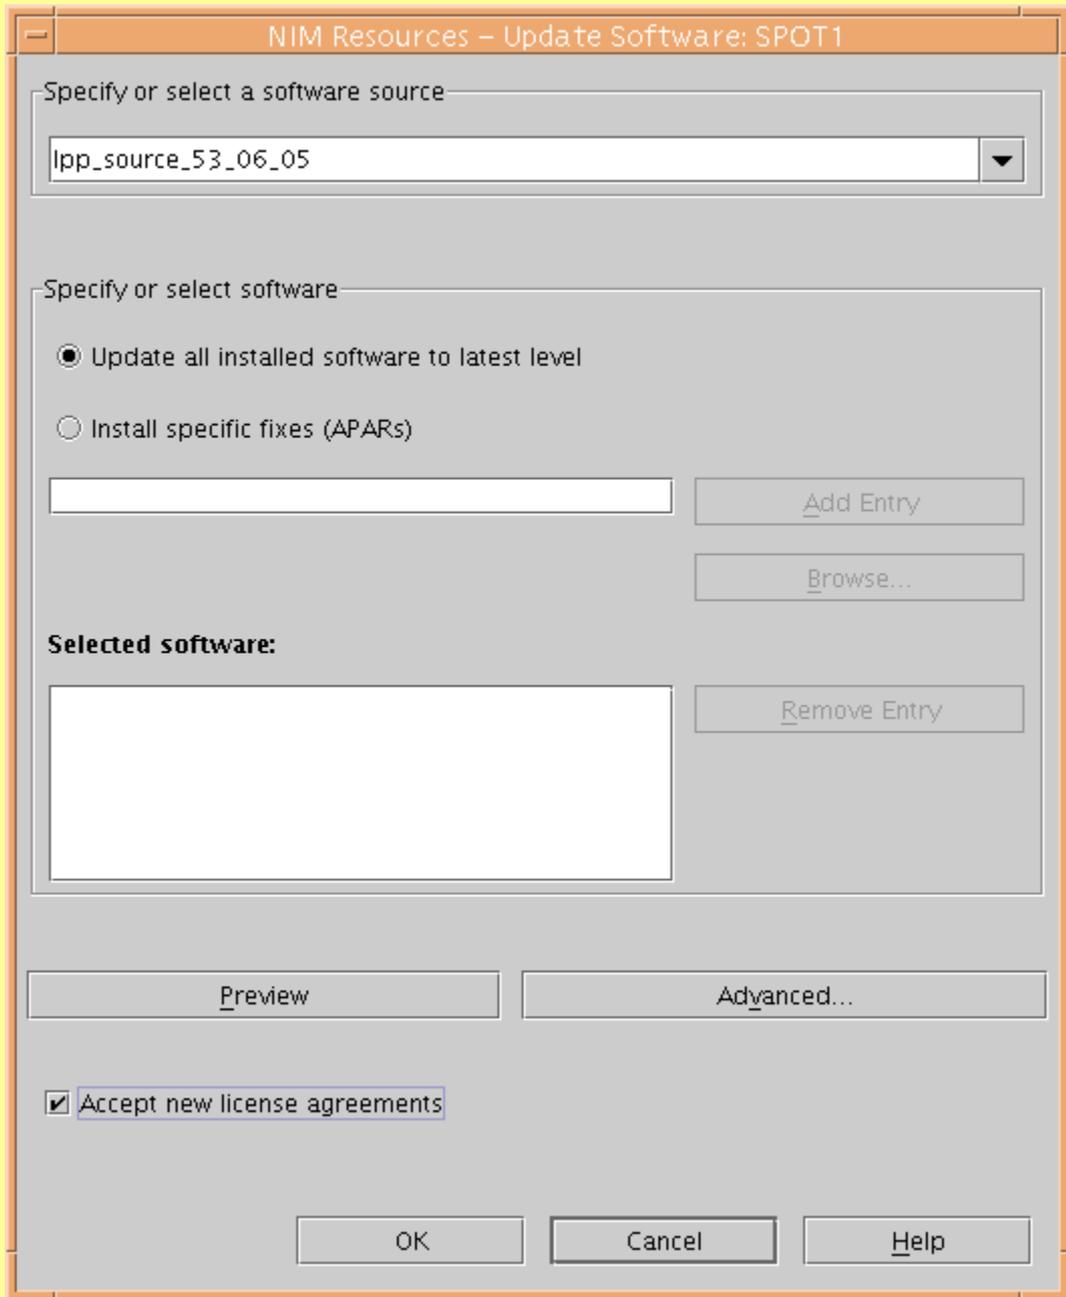


Next is to patch the SPOT which is with TL04 with the TL06 lpp and then with TL06_SP5 lpp.









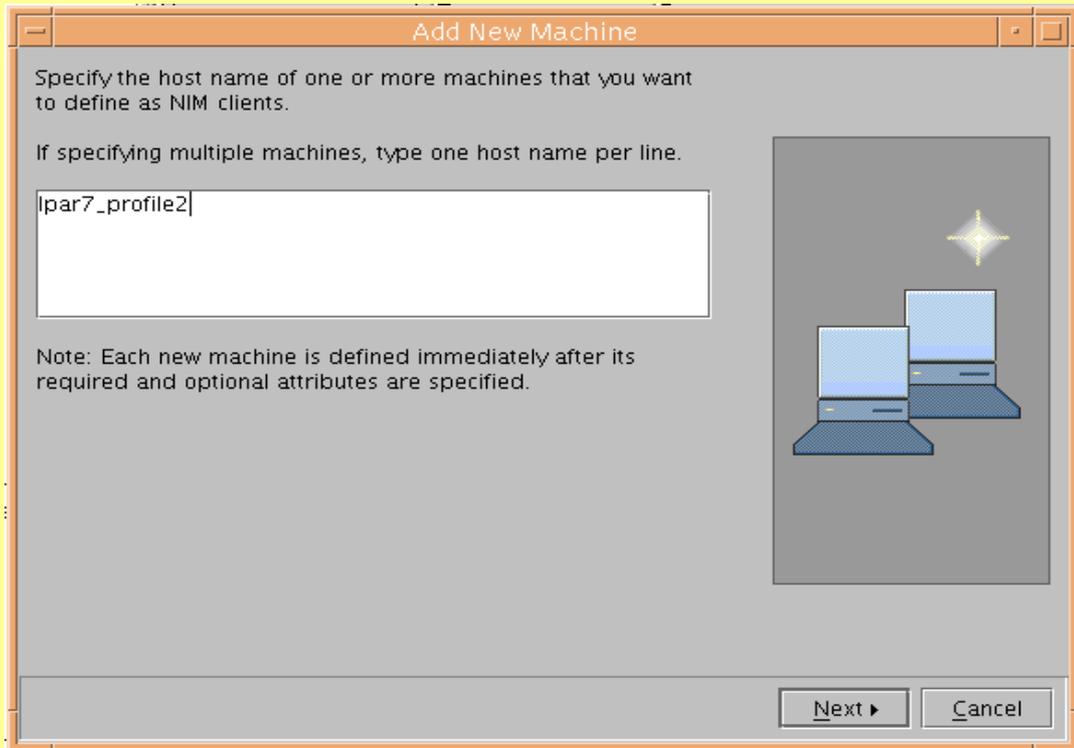


With the above steps, SPOT and LPPs are configured and ready to use.

NIM Client Configuration, Installation & Patching

Edit the /etc/hosts file and include name resolution for the new client lpar

Then in wsm, choose the machines option and invoke : create a new machine.



Add New Machine

To define the required NIM attributes of the host lpar7_profile2 using the default values shown, click Next.

To specify new values, modify the default values.

NIM name:

Machine type:

Hardware platform:

Cable type:

Communication Protocol used by client:

Network boot kernel:

Uniprocessor

Multiprocessor

Add New Machine

To define the optional NIM attributes of the host lpar7_profile2 using the default values shown, click Next.

To specify new values, modify the default values.

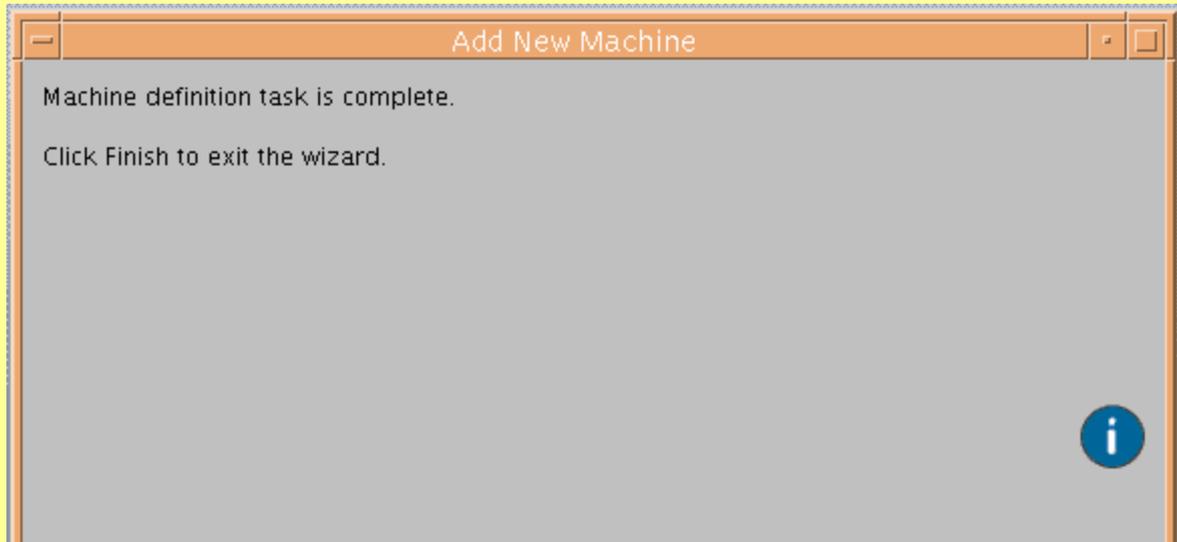
Network adapter logical device name:

IPL ROM emulation device:

Network adapter hardware address:

Machine group:

Comments:

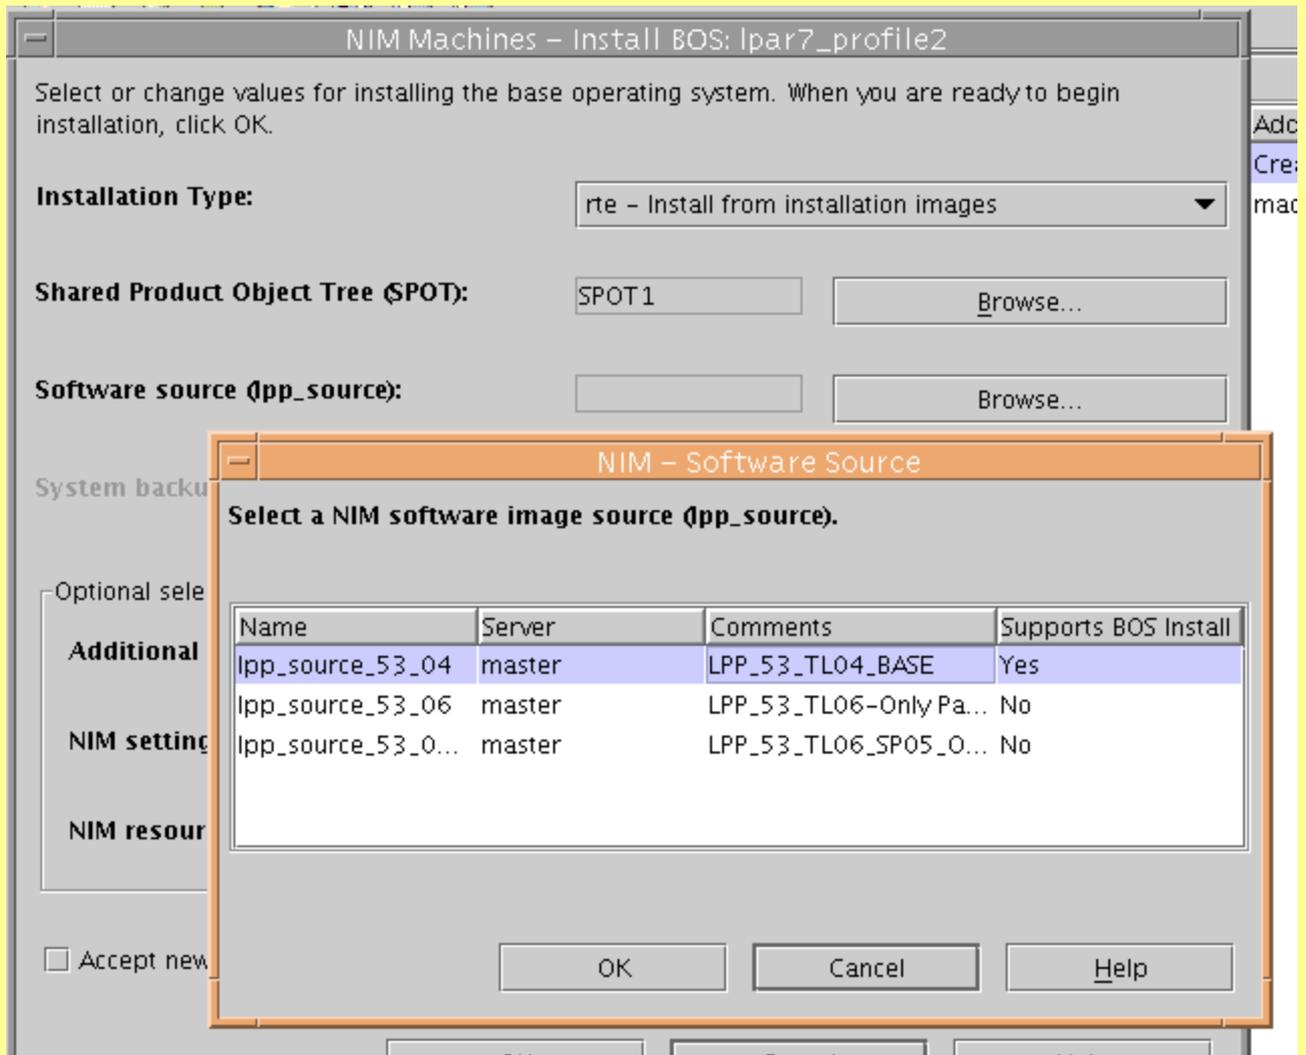
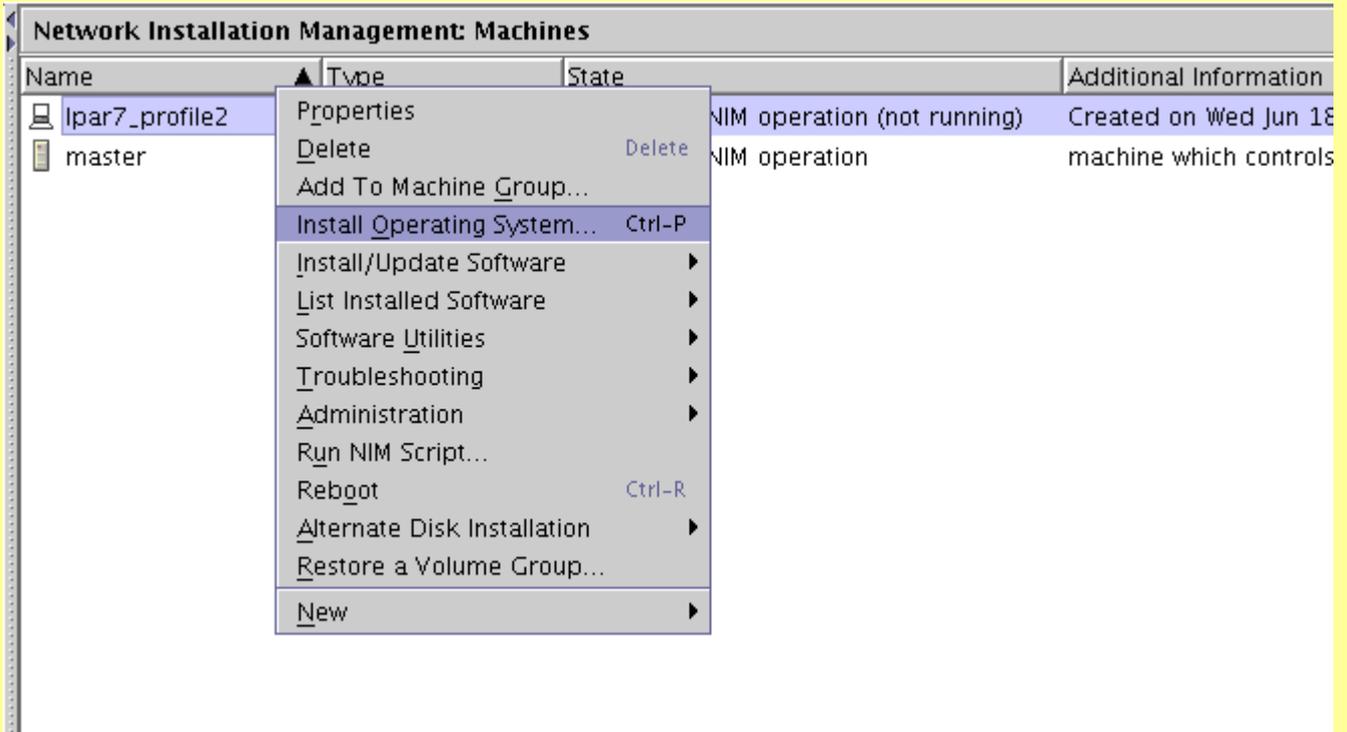


Network Installation Management: Machines

Name	Type	State	Additional Information
lpar7_profile2	standalone	ready for a NIM operation (not running)	Created on Wed Jun 18 19 25 41
master	master	ready for a NIM operation	machine which controls the NIM e

```
# Internet Address      Hostname      # Comments
# 192.9.200.1          net0sample   # ethernet name/address
# 128.100.0.1          token0sample  # token ring name/address
# 10.2.0.2             x25sample    # x.25 name/address
127.0.0.1              loopback localhost    # loopback (lo0) name/address
10.10.10.22           lpar1_profile
10.10.10.26           localhost
10.10.10.21           standbysvr localhost
10.10.10.27           lpar7_profile2
```

Now Resources are ready and client machine is ready for BOS deployment.



NIM Machines – Install BOS: lpar7_profile2

Select or change values for installing the base operating system. When you are ready to begin installation, click OK.

Installation Type: rte - Install from installation images

Shared Product Object Tree (SPOT): SPOT1

Software source (lpp_source): lpp_source_53_04

System backup image (mksysb):

Optional selections

Additional software to install:

NIM settings:

NIM resources:

Accept new license agreements

In the above given option for LPP resource, only one which is a BOS base package supports BOS installation.

Installing the Client machine.

```
PowerPC Firmware
Version SF240_338
SMS 1.6 (c) Copyright IBM Corp. 2000,2005 All rights reserved.
```

Main Menu

1. Select Language
2. Setup Remote IPL (Initial Program Load)
3. Change SCSI Settings
4. Select Console
5. Select Boot Options

```
PowerPC Firmware
Version SF240_338
SMS 1.6 (c) Copyright IBM Corp. 2000,2005 All rights reserved.
```

NIC Adapters

	Device	Location Code	Hardware Address
1.	Port 1 - IBM 2 PORT 10/100/100	U5791.001.99B0TG4-P1-C01-T1	00145ec7d6ce
2.	Port 2 - IBM 2 PORT 10/100/100	U5791.001.99B0TG4-P1-C01-T2	00145ec7d6cf

```
PowerPC Firmware
Version SF240_338
SMS 1.6 (c) Copyright IBM Corp. 2000,2005 All rights reserved.
```

Network Parameters

```
Port 1 - IBM 2 PORT 10/100/1000 Base-TX PCI-X Adapter: U5791.001.99B0TG4-P1-C01
1. IP Parameters
2. Adapter Configuration
3. Ping Test
4. Advanced Setup: BOOTP
```

```
PowerPC Firmware
Version SF240_338
SMS 1.6 (c) Copyright IBM Corp. 2000,2005 All rights reserved.
```

IP Parameters

```
Port 1 - IBM 2 PORT 10/100/1000 Base-TX PCI-X Adapter: U5791.001.99B0TG4-P1-C01
1. Client IP Address [10.10.10.27]
2. Server IP Address [10.10.10.21]
3. Gateway IP Address [000.000.000.000]
4. Subnet Mask [255.255.255.000]
```

Client IP : IP address for the lpar, Server IP : IP address of the nim_master.

```
PowerPC Firmware
Version SF240_338
SMS 1.6 (c) Copyright IBM Corp. 2000,2005 All rights reserved.
```

Multiboot

1. Select Install/Boot Device
2. Configure Boot Device Order
3. Multiboot Startup <OFF>

```
PowerPC Firmware
Version SF240_338
SMS 1.6 (c) Copyright IBM Corp. 2000,2005 All rights reserved.
```

Adapter Configuration

```
Port 1 - IBM 2 PORT 10/100/1000 Base-TX PCI-X Adapter: U5791.001.99B0TG4-P1-C01
1. Speed,Duplex
2. Spanning Tree Enabled
3. Protocol
```

```
PowerPC Firmware
Version SF240_338
SMS 1.6 (c) Copyright IBM Corp. 2000,2005 All rights reserved.
```

```
Spanning Tree Enabled
Port 1 - IBM 2 PORT 10/100/1000 Base-TX PCI-X Adapter: U5791.001.99B0TG4-P1-C01
1.  Yes
2.  No  <===
```

```
PowerPC Firmware
Version SF240_338
SMS 1.6 (c) Copyright IBM Corp. 2000,2005 All rights reserved.
```

```
Select Device
Device  Current  Device
Number  Position  Name
1.      1      Ethernet
          ( loc=U5791.001.99B0TG4-P1-C01-T1 )
2.      -      Ethernet
          ( loc=U5791.001.99B0TG4-P1-C01-T2 )
```

```
Version SF240_338
SMS 1.6 (c) Copyright IBM Corp. 2000,2005 All rights reserved.
```

```
Select Task

Ethernet
  ( loc=U5791.001.99B0TG4-P1-C01-T1 )

1.  Information
2.  Normal Mode Boot
3.  Service Mode Boot
```

```

BOOTP: chosen-network-type = ethernet,auto,rj45,auto
BOOTP: server   IP =      10.10.10.21
BOOTP: requested filename =
BOOTP: client   IP =      10.10.10.27
BOOTP: client   HW addr =  0 14 5e c7 d6 ce
BOOTP: gateway  IP =      0.0.0.0
BOOTP: device   /pci@800000020000017/pci@2/ethernet@1
BOOTP: loc-code U5791.001.99B0TG4-P1-C01-T1

```

```

BOOTP R = 1 BOOTP S = 2
FILE: /tftpboot/lpar7_profile2

```

```

BOOTP R = 1 BOOTP S = 2
FILE: /tftpboot/lpar7_profile2
FINAL Packet Count = 25834
FINAL File Size = 13226496 bytes.
load-base=0x4000
real-base=0x2000000

```

Elapsed time since release of system processors: 3415 mins 1 secs

```

-----
                        Welcome to AIX.
                        boot image timestamp: 14:01 06/18
                        The current time and date: 14:42:17 06/18/2008
                        number of processors: 1      size of memory: 256MB
boot device: /pci@800000020000017/pci@2/ethernet@1:10.10.10.21,,10.10.10.27,000.
000.000.000,00,00
                        kernel size: 12105402; 32 bit kernel
-----

```

```
***** Please define the System Console. *****
```

```
Type a 1 and press Enter to use this terminal as the  
system console.
```

```
Pour definir ce terminal comme console systeme, appuyez  
sur 1 puis sur Entree.
```

```
Taste 1 und anschliessend die Eingabetaste druecken, um  
diese Datenstation als Systemkonsole zu verwenden.
```

```
Premere il tasto 1 ed Invio per usare questo terminal  
come console.
```

```
Escriba 1 y pulse Intro para utilizar esta terminal como  
consola del sistema.
```

```
Escriviu 1 i premeu Intro per utilitzar aquest  
terminal com a consola del sistema.
```

```
Digite um 1 e pressione Enter para utilizar este terminal  
como console do sistema.
```

```
>>> 1 Type 1 and press Enter to have English during install.
```

```
88 Help ?
```

```
>>> Choice [1]:
```

Error Warning

There are no disks available on this system.

To reboot the system, press reset.

>>> 1 Continue with Install

Welcome to Base Operating System
Installation and Maintenance

Type the number of your choice and press Enter. Choice is indicated by >>>.

>>> 1 Start Install Now with Default Settings

2 Change/Show Installation Settings and Install

3 Start Maintenance Mode for System Recovery

Install Options

1. Enable Trusted Computing Base..... no
2. Enable CAPP and EAL4+ Technology..... no
(English only, 64-bit kernel enablement, JFS2 file systems)
3. Enable 64-bit Kernel..... yes
4. Create JFS2 File Systems..... yes
5. Graphics Software..... yes
6. Enable System Backups to install any system..... yes
(Installs all devices and kernels)

Installation and Settings

Either type 0 and press Enter to install with current settings, or type the number of the setting you want to change and press Enter.

- 1 System Settings:
 - Method of Installation.....New and Complete Overwrite
 - Disk Where You Want to Install.....hdisk0

- 2 Primary Language Environment Settings (AFTER Install):
 - Cultural Convention.....C (POSIX)
 - Language.....C (POSIX)
 - Keyboard.....C (POSIX)

3 More Options (Desktop, Security, Kernel, Software, ...)

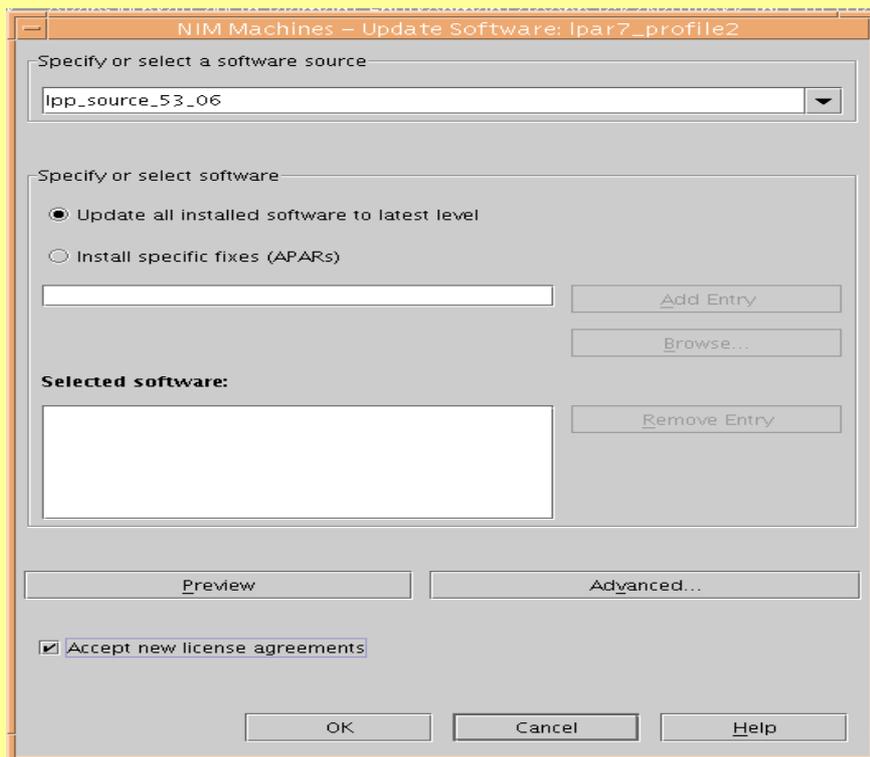
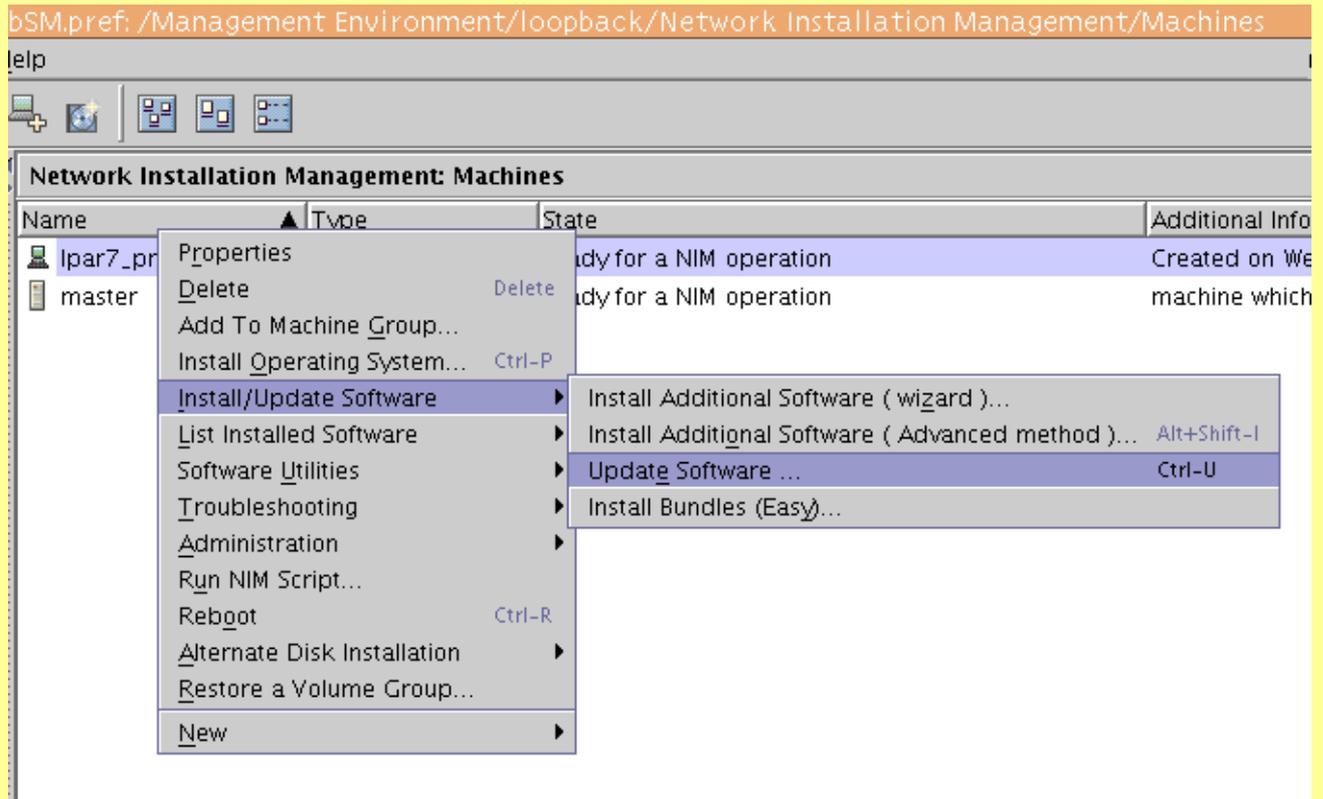
>>> 0 Install with the settings listed above.

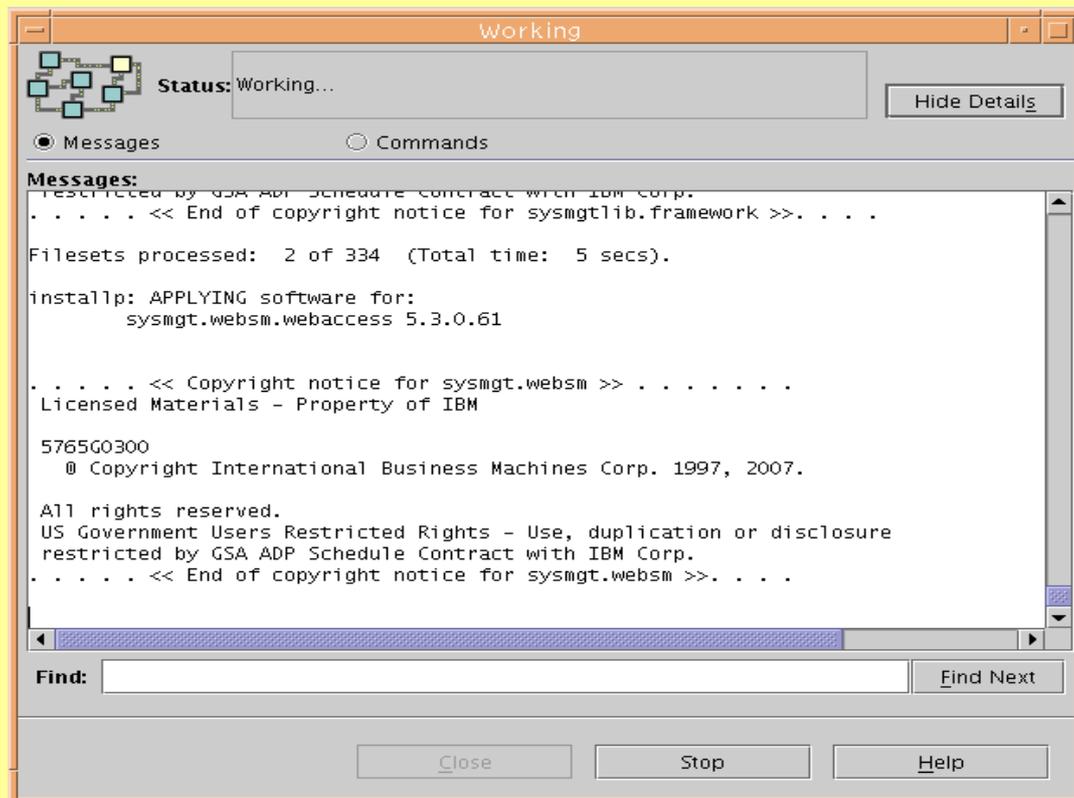
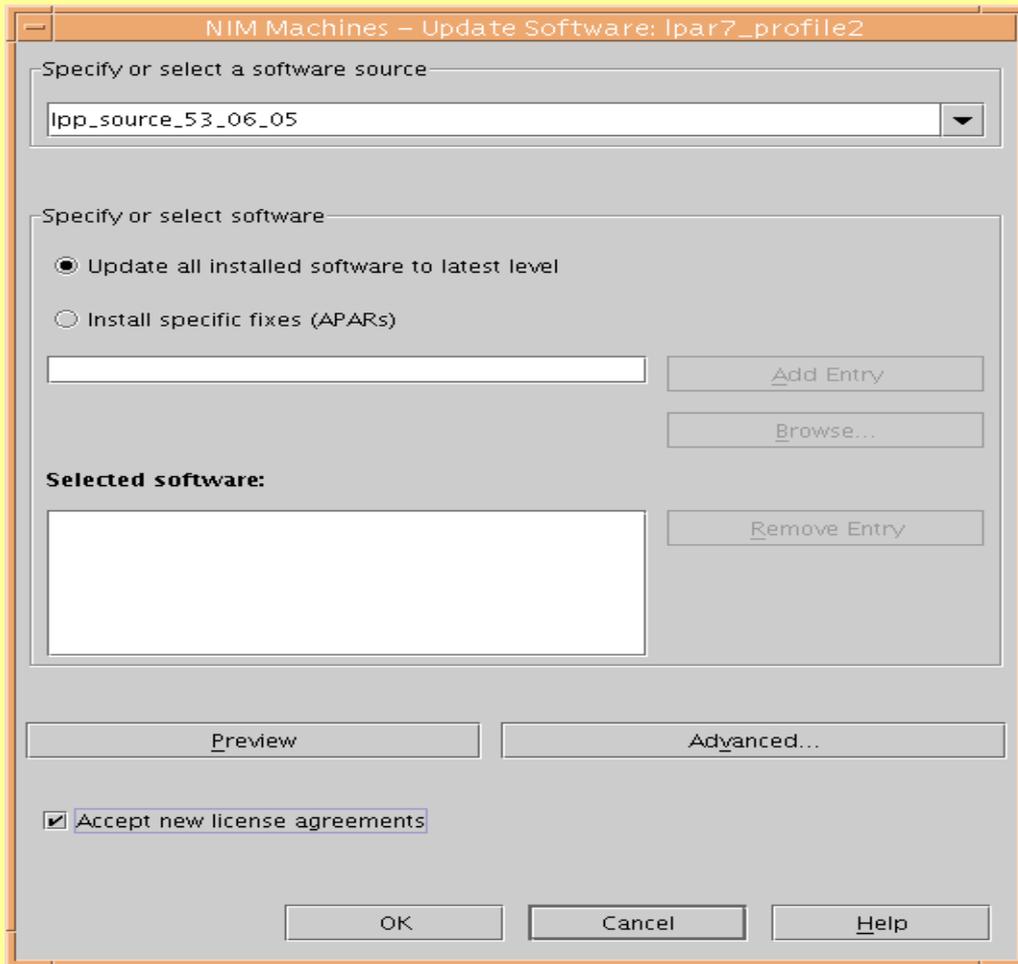
Network Installation Management: Machines		
Machine Name	State	Additional Information
hdalone	customization is being performed (in the process of...	BOS install 7% complete : Restoring base operating .
nter	ready for a NIM operation	machine which controls the NIM environment

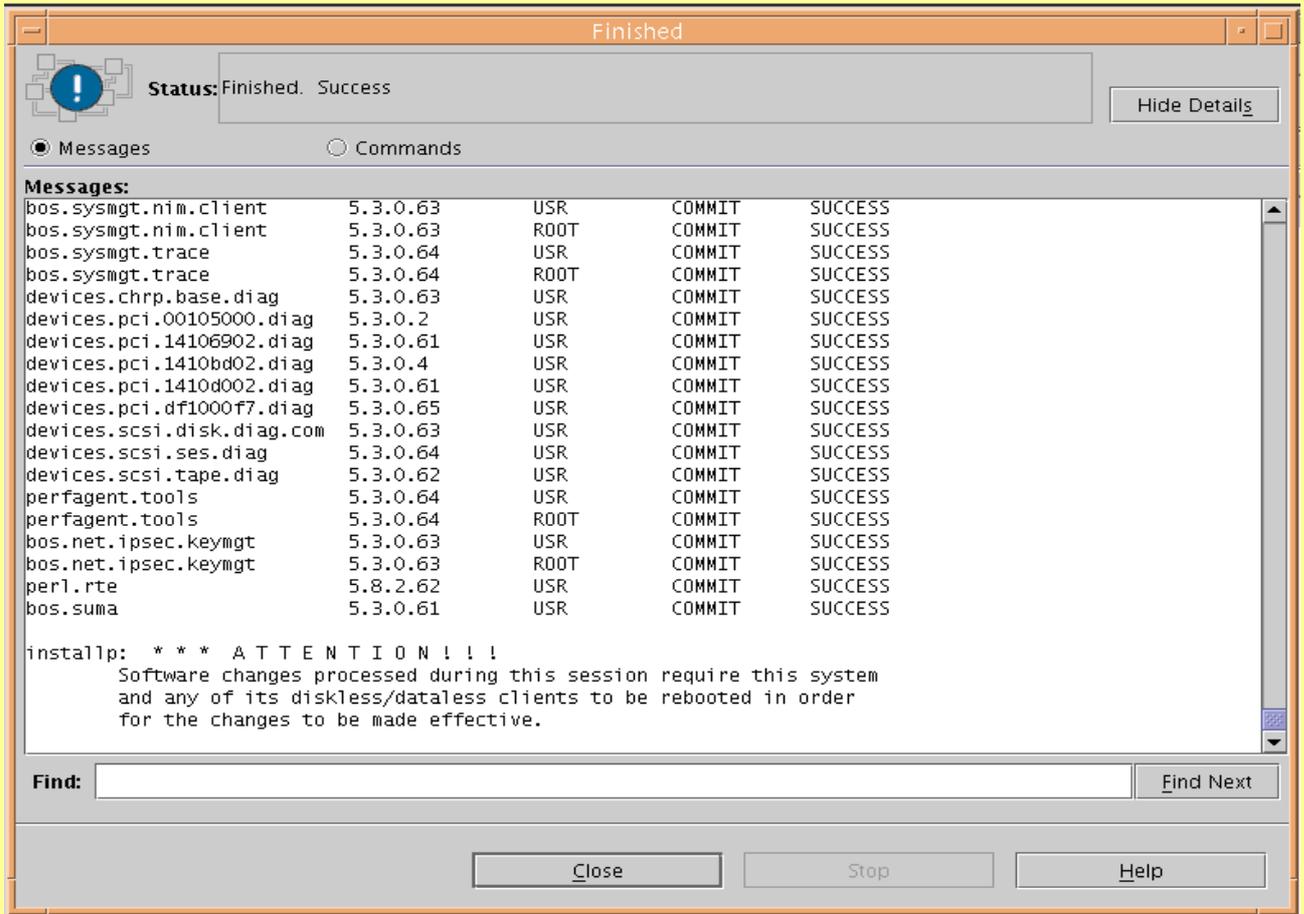
This continues and installation finishes. NIM installation from the lpp is exactly the same as the installation from a installation CD, but with the difference that it is through network. Multiple clients can be installed simultaneously using NIM. On a 1Gbps network, it takes around 10-15minutes for one deployment.

It is only possible to install the BOS from a lpp made out of the BASE package.
Thus the current deployed AIX is version 5.3 TL04.

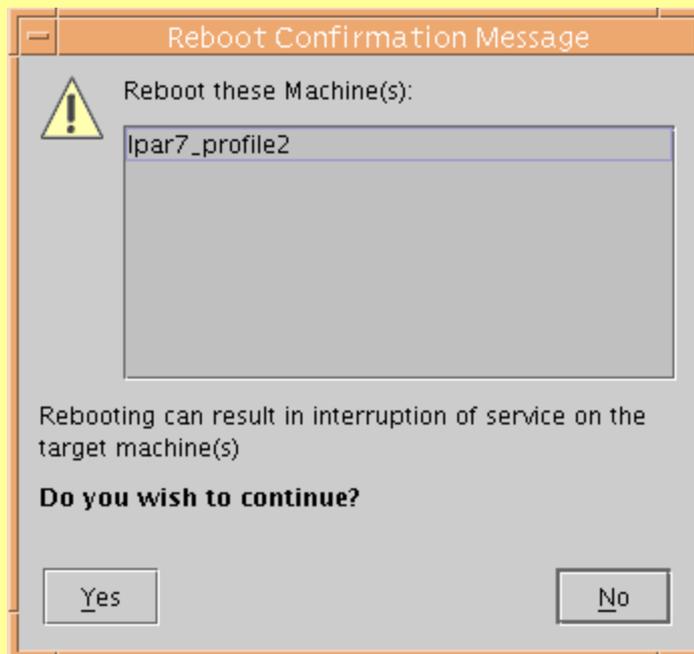
Now its time to patch the client to TL06+SP5



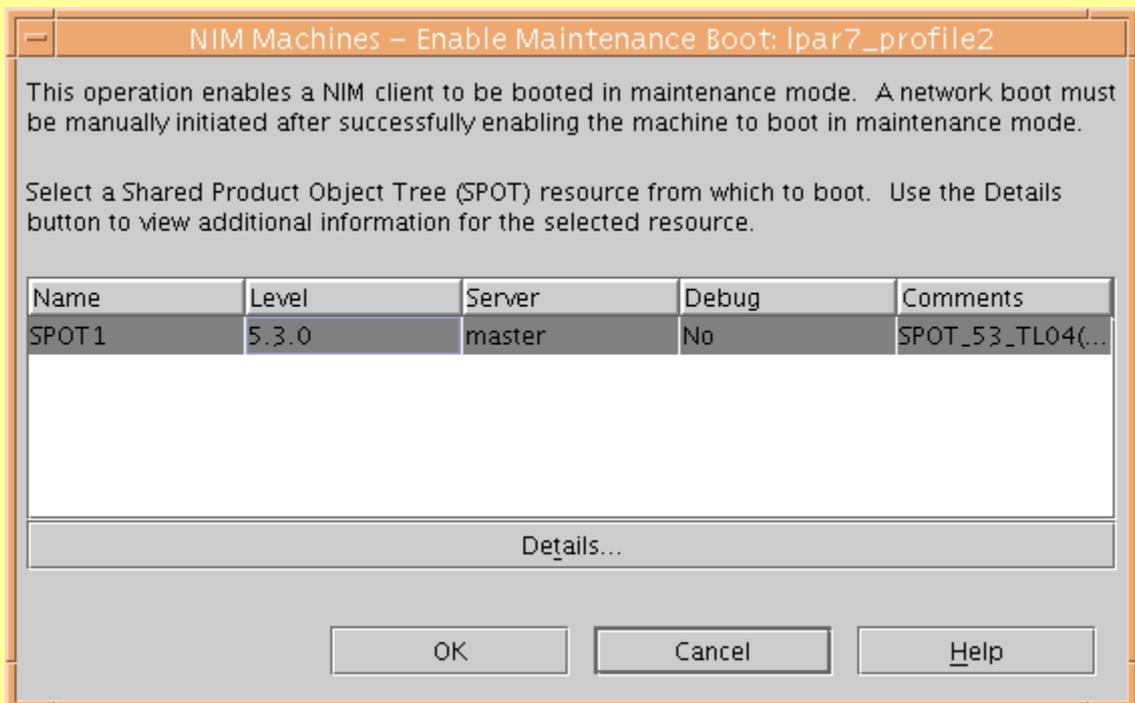
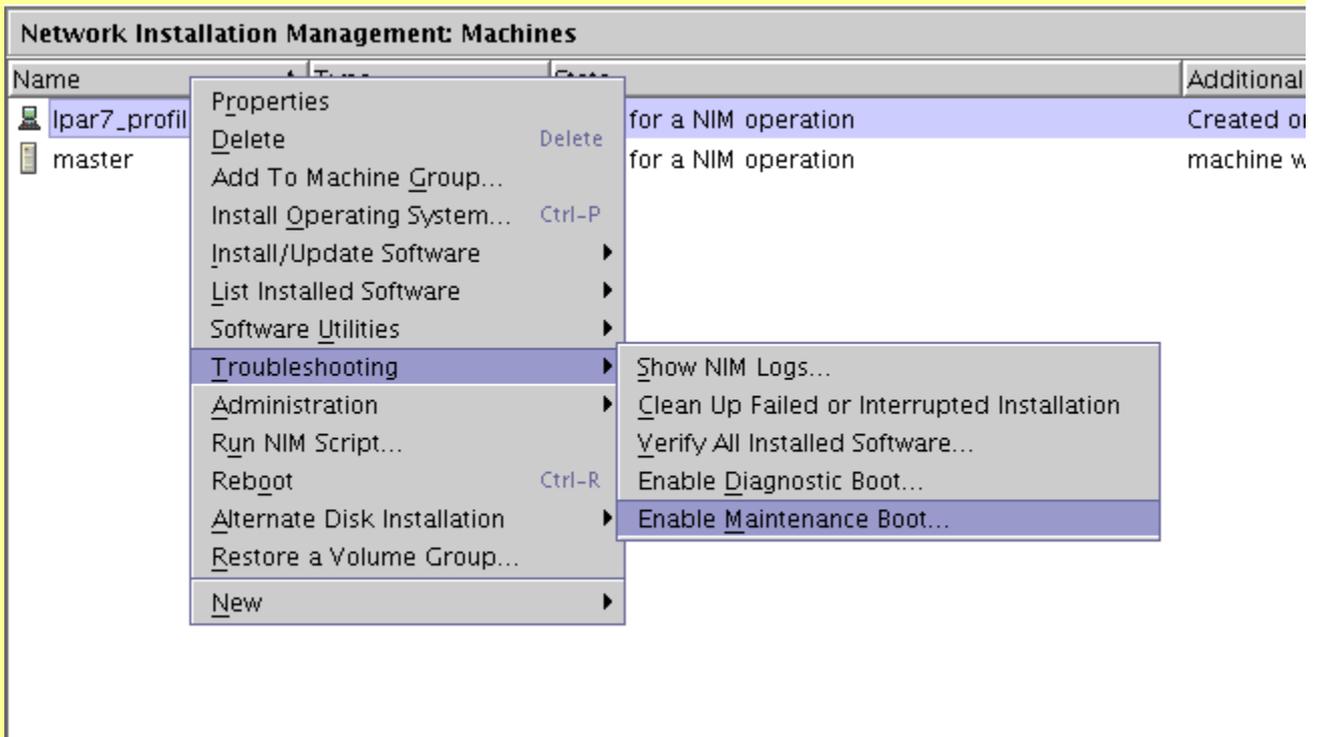




You can reboot the nim client from the nim_master itself.



Enabling Maintenance mode boot with nim_master.



Network Installation Management: Machines			
Name	Type	State	Additional Information
lpar7_profile2	standalone	maintenance boot has been enabled	Created on Wed Jun 1
master	master	ready for a NIM operation	machine which control

Boot up the client machine using network card ...

```
PowerPC Firmware
Version SF240_338
SMS 1.6 (c) Copyright IBM Corp. 2000,2005 All rights reserved.
```

```
-----
Select Device
```

```
Device   Current   Device
Number  Position  Name
1.      -         Ethernet
          ( loc=U5791.001.99B0TG4-P1-C01-T1 )
2.      -         Ethernet
          ( loc=U5791.001.99B0TG4-P1-C01-T2 )
```

```
Select Task
```

```
Ethernet
  ( loc=U5791.001.99B0TG4-P1-C01-T1 )

1.   Information
2.   Normal Mode Boot
3.   Service Mode Boot
```

```
BOOTP R = 1 BOOTP S = 2
FILE: /tftpboot/lpar7_profile2
FINAL Packet Count = 25834
FINAL File Size = 13226496 bytes.
load-base=0x4000
real-base=0x2000000
```

```
Elapsed time since release of system processors: 3563 mins 57 secs
```

```
-----
                        Welcome to AIX.
                boot image timestamp: 14:01 06/18
                The current time and date: 17:10:35 06/18/2008
                number of processors: 1    size of memory: 512MB
boot device: /pci@800000020000017/pci@2/ethernet@1:10.10.10.21,,10.10.10.27,000.
000.000.000,00,00
                kernel size: 12105402; 32 bit kernel
-----
```

Maintenance

Type the number of your choice and press Enter.

- >>> 1 Access a Root Volume Group
- 2 Copy a System Dump to Removable Media
- 3 Access Advanced Maintenance Functions
- 4 Erase Disks

Volume Group Information

Volume Group ID 00c28f2000004c000000011a9c2c1c59 includes the following
logical volumes:

hd5	hd6	hd8	hd4	hd2	hd9var
hd3	hd1	hd10opt			

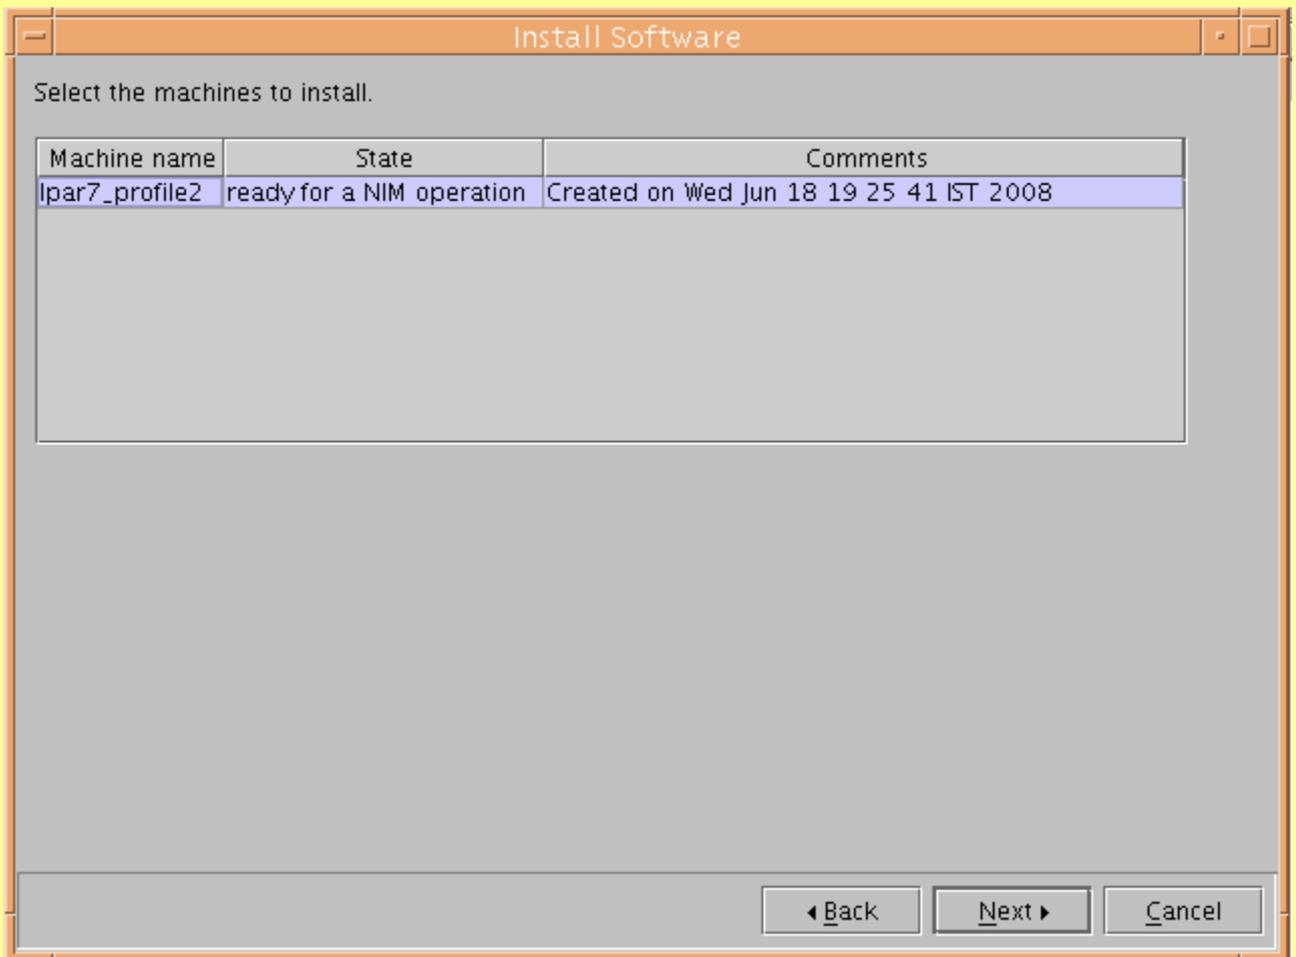
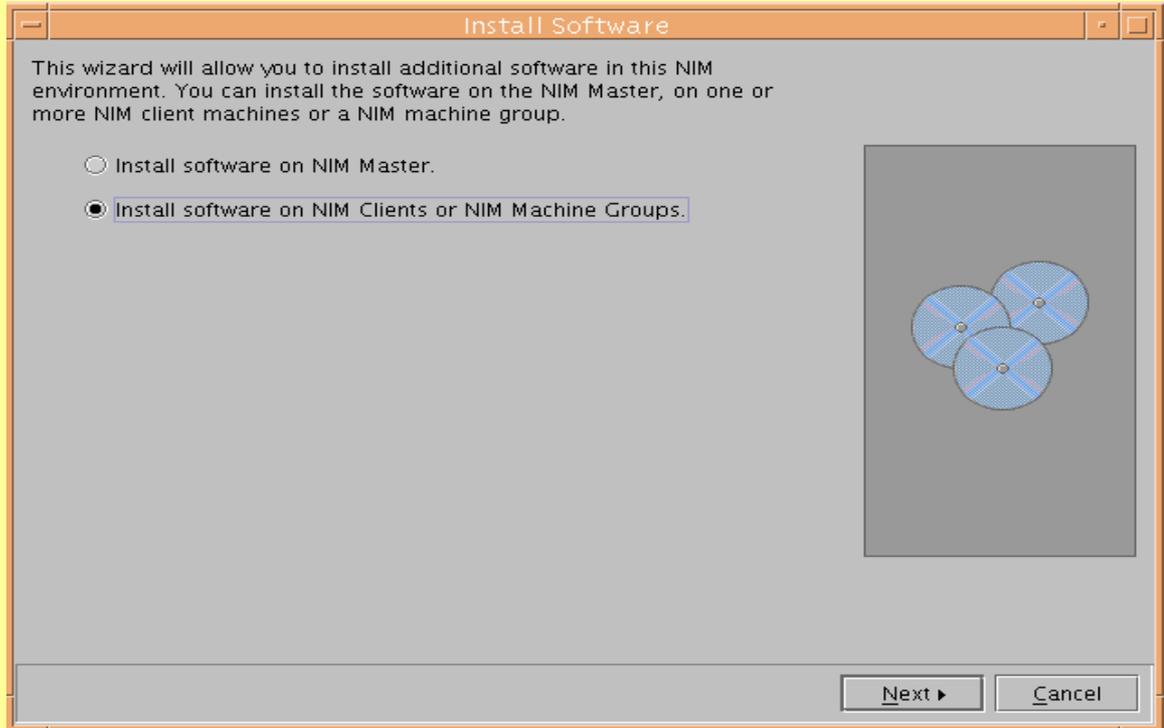
Type the number of your choice and press Enter.

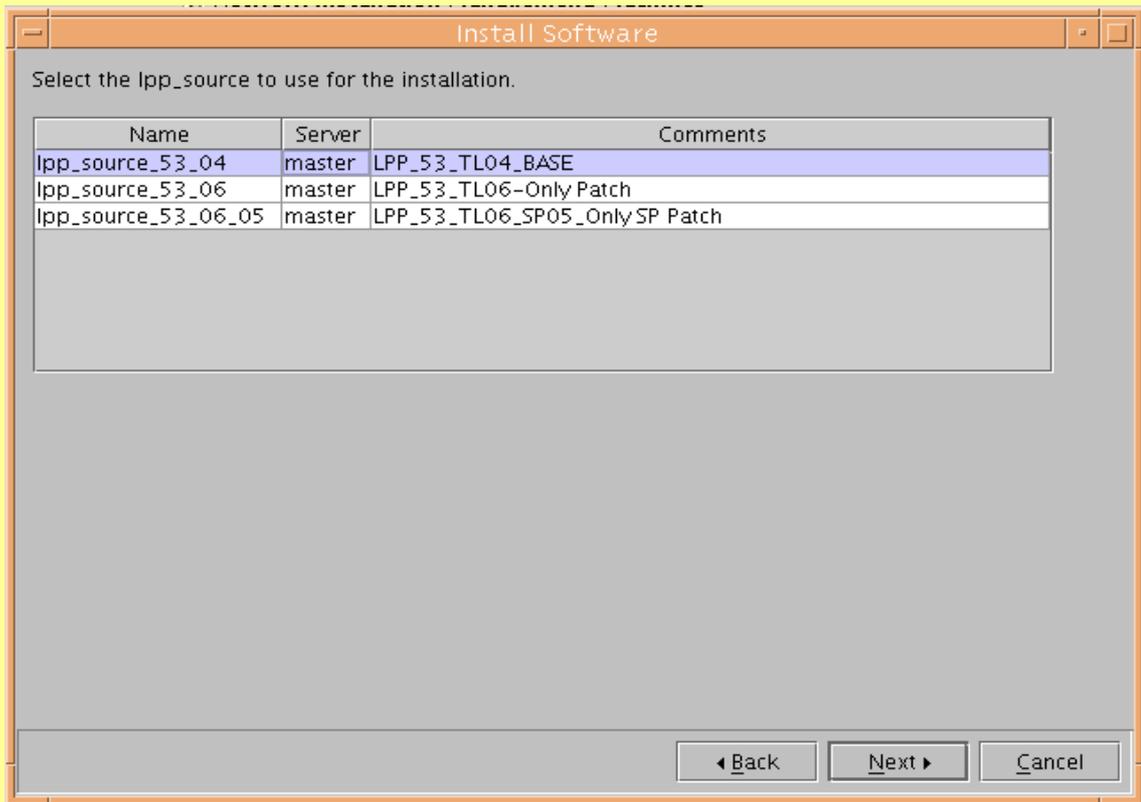
- 1) Access this Volume Group and start a shell
- 2) Access this Volume Group and start a shell before mounting filesystems

99) Previous Menu

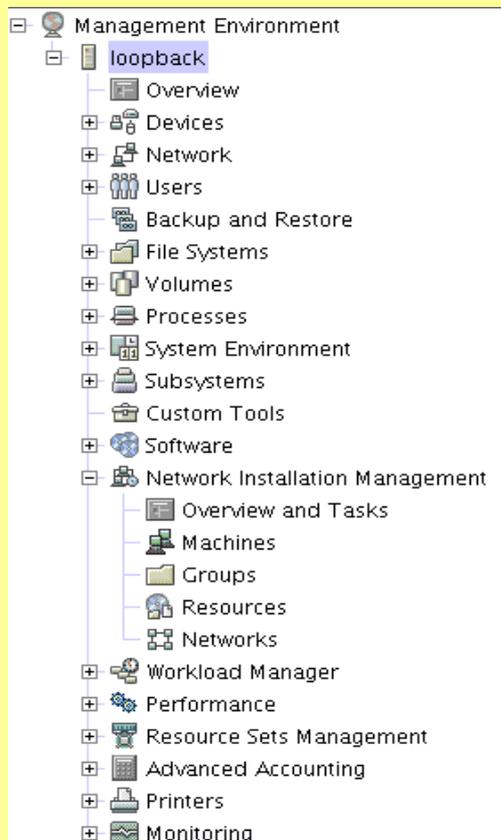
Choice [99]: 1

Installing software using NIM.





[NIM restoration of the mksysb image.](#)



Network Installation Management: Resources				
Name	Type	State	Additional Information	
 lpp_source_53_04	lpp_source	ready for use	LPP_53_TL04_BASE	
 lpp_source_53_06	lpp_source	ready for use	LPP_53_TL06-Only Patch	
 lpp_source_53_06_05	lpp_source	ready for use	LPP_53_TL06_SP05_Only SP Patch	
 SPOT1	spot	ready for use	SPOT_53_TL04(base)+TL06+SP05	

Add New Resource

Select the type of BOS Installation resource to add to the NIM environment.

Required resources:

- lpp_source - Directory of installation images
- SPOT - Provides network boot & installation support
- lpp_source & SPOT - Resource pair (2) created from same input images
- mksysb - System backup image in "mksysb" format

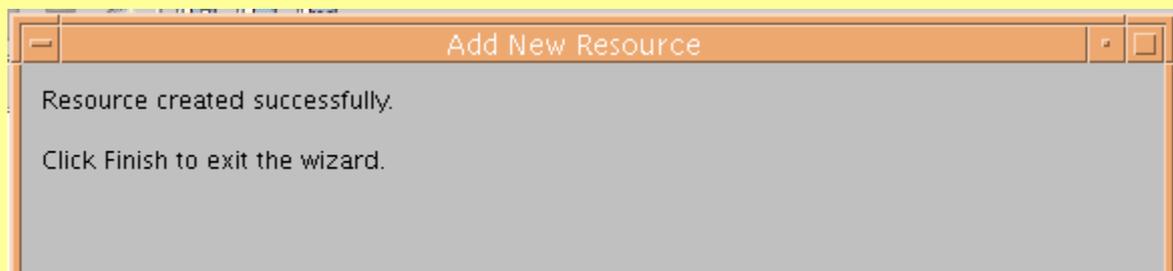
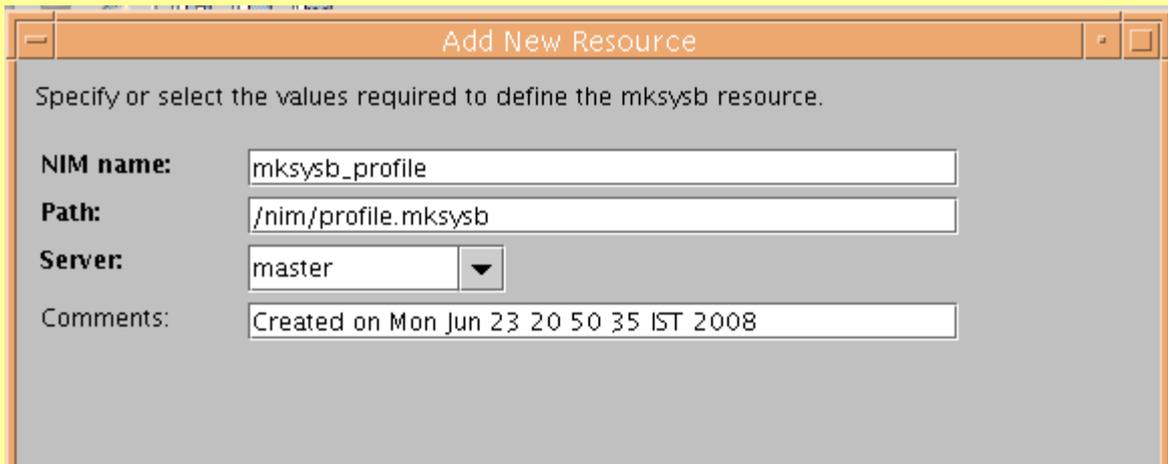
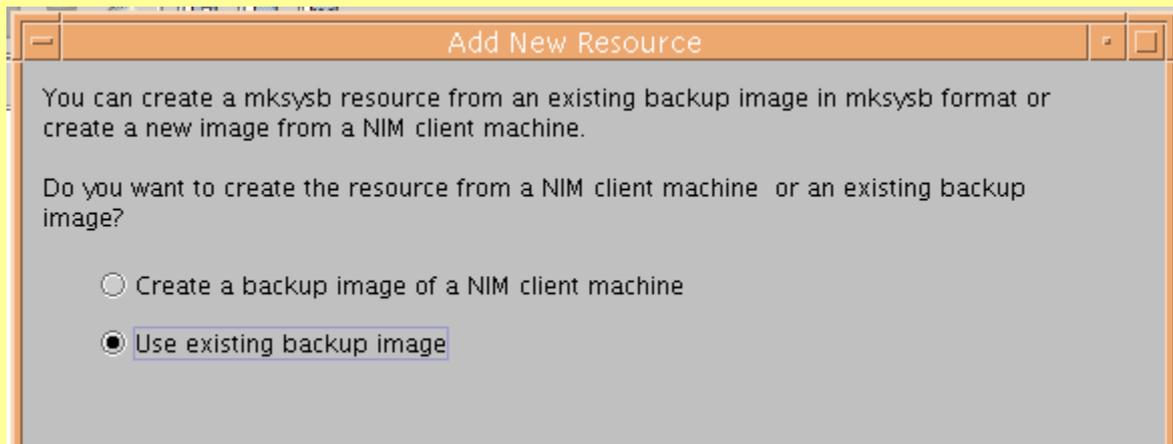
Optional resources:

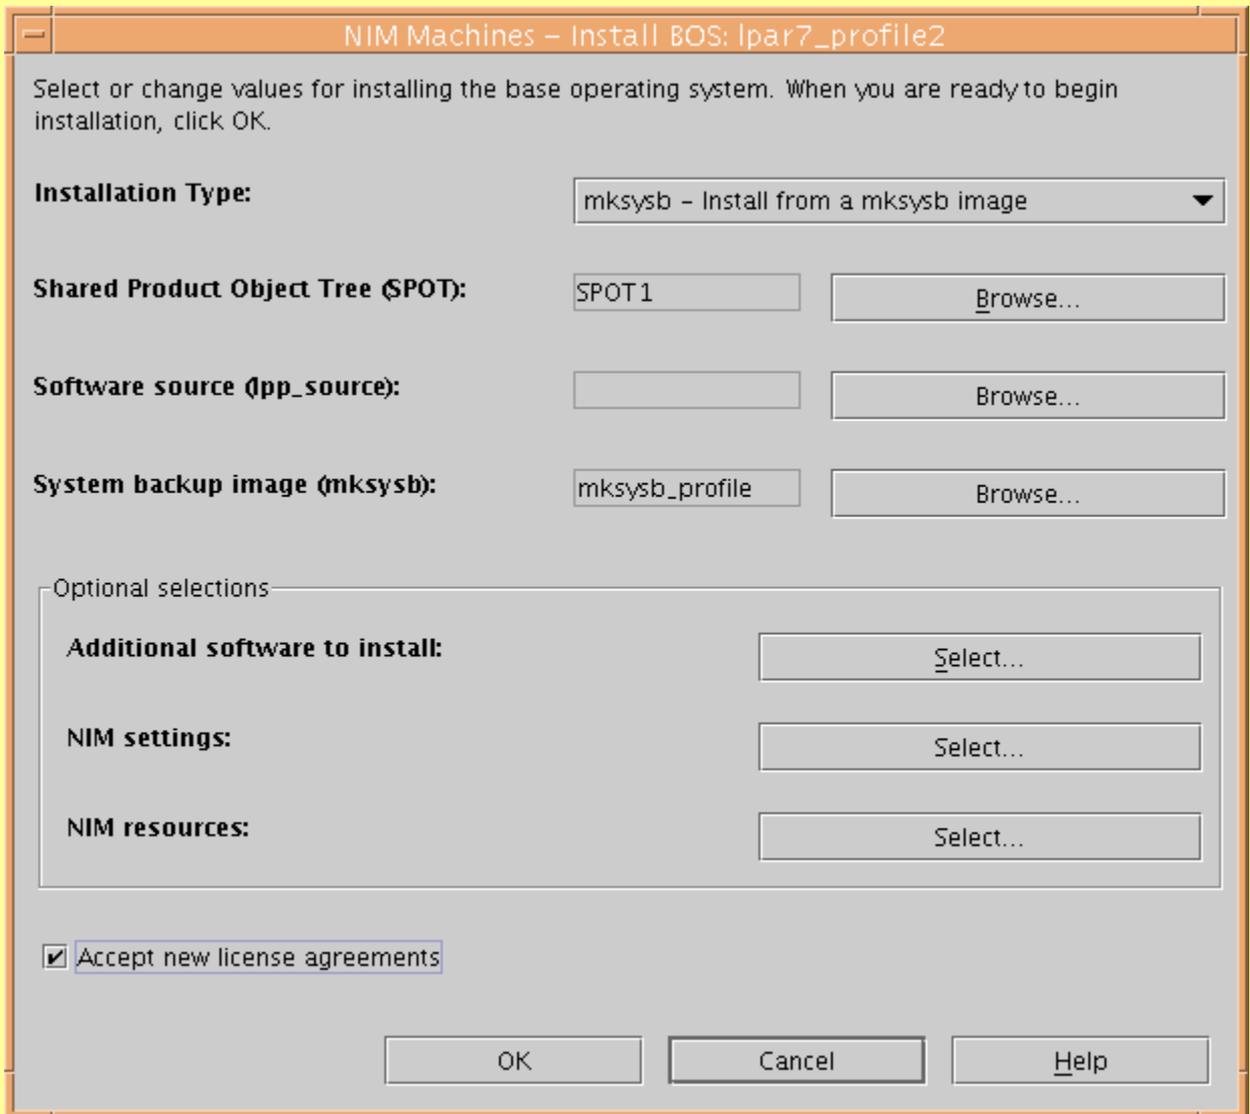
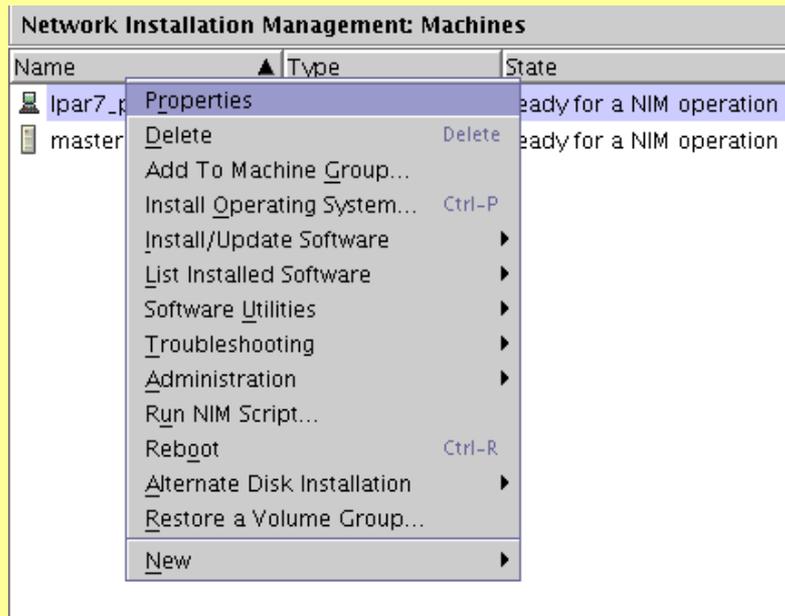
- bosinst_data - Customizes BOS install flow (automates choices)
- installp_bundle - Lists additional software to install after BOS
- script - Customization program executed after installation
- fb_script - Executable script added to the /etc/firstboot file, and is run at first reboot after bos install to configure devices
- resolv_conf - Provides DNS entries for /etc/resolv.conf file

Add New Resource

All necessary resource characteristics have been specified. To add the mksysb resource now, click Next.

To verify settings before defining the resource, click View Settings.





Network Installation Management: Machines		
Name	Type	State
lpar7_profile2	standalone	BOS installation has been enabled (not running)
master	master	ready for a NIM operation

```
PowerPC Firmware
Version SF240_338
SMS 1.6 (c) Copyright IBM Corp. 2000,2005 All rights reserved.
```

```
-----
Device Information
IBM,FW-REVISION-LEVEL: 210
SPEED,DUPLEX      : auto,auto
CLIENT IP ADDRESS : 10.10.10.27
SERVER IP ADDRESS : 10.10.10.21
GATEWAY IP ADDRESS : 000.000.000.000
SUBNET MASK       : 255.255.255.000
ETHERNET PROTOCOL : Standard
SPANNING TREE ENABLED: No
```

```
Parent Information
NAME           : pci
DEVICE-TYPE    : pci
```

```
Version SF240_338
SMS 1.6 (c) Copyright IBM Corp. 2000,2005 All rights reserved.
```

```
-----
Device Information
IBM,FW-ADAPTER-NAME: Port 1 - IBM 2 PORT 10/100/1000 Base-TX PCI-X Adapter
                    /pci@8000000020000017/pci@2/ethernet@1
                    : (Bootable)
DEVICE              : Ethernet
                    ( loc=U5791.001.99B0TG4-P1-C01-T1 )
NAME                : ethernet
DEVICE-TYPE         : network
SUPPORTED-NETWORK-TYPES:
                    : ethernet,auto,rj45,auto   <==  chosen
                    : ethernet,10,rj45,half
                    : ethernet,10,rj45,full
                    : ethernet,100,rj45,half
                    : ethernet,100,rj45,full
                    : ethernet,1000,rj45,full
MAC-ADDRESS         : 00145ec7d6ce
```

```
BOOTP: chosen-network-type = ethernet,auto,rj45,auto
BOOTP: server IP = 10.10.10.21
BOOTP: requested filename =
BOOTP: client IP = 10.10.10.27
BOOTP: client HW addr = 0 14 5e c7 d6 ce
BOOTP: gateway IP = 0.0.0.0
BOOTP: device /pci@800000020000017/pci@2/ethernet@1
BOOTP: loc-code U5791.001.99B0TG4-P1-C01-T1
```

```
BOOTP R = 1 BOOTP S = 2
FILE: /tftpboot/lpar7_profile2
FINAL Packet Count = 25834
FINAL File Size = 13226496 bytes.
load-base=0x4000
real-base=0x2000000
```

```
Elapsed time since release of system processors: 10679 mins 37 secs
```

```
-----
Welcome to AIX.
boot image timestamp: 14:01 06/18
The current time and date: 15:46:20 06/23/2008
number of processors: 1 size of memory: 512MB
boot device: /pci@800000020000017/pci@2/ethernet@1:10.10.10.21,,10.10.10.27,
000.000.000,00,00
kernel size: 12105402; 32 bit kernel
-----
```

Network Installation Management: Machines		
Name	Type	State
lpar7_profile2	standalone	BOS installation has been enabled (in the process of booting)
master	master	ready for a NIM operation

Change Disk(s) Where You Want to Install

Type one or more numbers for the disk(s) to be used for installation and press Enter. To cancel a choice, type the corresponding number and Press Enter. At least one bootable disk must be selected. The current choice is indicated by >>>.

	Name	Location Code	Size(MB)	VG Status	Bootable	Maps
>>> 1	hdisk0	02-08-00-8,0	140013	rootvg	Yes	No
>>> 2	hdisk1	02-08-00-9,0	140013	none	Yes	No

Installing Base Operating System

Please wait...

Approximate % tasks complete	Elapsed time (in minutes)
12	3

7% of mksysb data restored.

```
Approximate      Elapsed time
% tasks complete (in minutes)

23              7      21% of mksysb data restored.
```

```
Approximate      Elapsed time
% tasks complete (in minutes)

39              13     43% of mksysb data restored.
```

```
Approximate      Elapsed time
% tasks complete (in minutes)

61              18     73% of mksysb data restored.
```

```
Approximate      Elapsed time
% tasks complete (in minutes)

72              22     87% of mksysb data restored.
```

```

Approximate      Elapsed time
% tasks complete (in minutes)

      82              23      Initializing disk environment.

```

```

Approximate      Elapsed time
% tasks complete (in minutes)

      86              25      Copying Cu* to disk.

```

```

Approximate      Elapsed time
% tasks complete (in minutes)

      88              29      Creating boot image.

```

System reboots after this step and completes the installation.

NIM is a very useful tool, but remember not to keep any production machine as nim_master.

Moreover a nim client has .rhosts entries allowing nim_master to access the client without password. This may be considered as a security violation in some of the deployments. Hence consider removing these .rhosts files after the nim deployments.



[rpm - Red Hat Package Manager](#)

RPM Package Manager (originally **Red Hat Package Manager**, abbreviated **RPM**) is a package management system. The name RPM refers to two things: a software package file format, and software packaged in this format. RPM was intended primarily for **Linux distributions**; the file format RPM is the baseline package format of the **Linux Standard Base**.

Originally developed by **Red Hat** for Red Hat Linux, RPM is now used by many **Linux distributions**. It has also been ported to some other operating systems, such as **Novell NetWare** (as of version 6.5 SP3) and **IBM's AIX** as of version 5.

"RPM" as it is used today is an example of a recursive initialism.

rpm is a powerful package manager, which can be used to build, install, query, verify, update, and uninstall individual software packages. A package consists of an archive of files, and package information, including name, version, and description.

One of the following basic modes must be selected:

Initialize Database,
 Rebuild Database,
 Build Package,
 Recompile Package,
 Build Package from Tarball,
 Query,
 Show Querytags,
 Install,
 Freshen,
 Uninstall,
 Verify, Signature Check, Resign,
 Add Signature, set owners and groups and Show Configuration.

RPM is ***free software***, released under the ***GNU GPL***.

In AIX you can call the man page of rpm using the command : ***man -M /opt/freeware/man/rpm***

Query the rpm database

rpm -q invokes the query mode of the command.

These are some of the useful rpm query options.

```
bash-3.00# rpm -qa
cdrecord-1.9-7
mkisofs-1.13-4
zip-2.3-3
bash-3.0-1
bash-doc-3.0-1
AIX-rpm-5.3.0.60-2
vnc-3.3.3r2-6
```

```
bash-3.00# rpm -ql zip-2.3-3
/opt/freeware/bin/zip
/opt/freeware/bin/zipcloak
/opt/freeware/bin/zipnote
/opt/freeware/bin/zipsplit
/opt/freeware/doc/zip-2.3
/opt/freeware/doc/zip-2.3/BUGS
/opt/freeware/doc/zip-2.3/CHANGES
/opt/freeware/doc/zip-2.3/LICENSE
/opt/freeware/doc/zip-2.3/MANUAL
/opt/freeware/doc/zip-2.3/README
/opt/freeware/doc/zip-2.3/TODO
/opt/freeware/doc/zip-2.3/WHATSNEW
/opt/freeware/doc/zip-2.3/WHERE
/opt/freeware/doc/zip-2.3/algorithm.txt
/opt/freeware/man/man1/zip.1
/usr/bin/zip
/usr/bin/zipcloak
/usr/bin/zipnote
/usr/bin/zipsplit
```

```
bash-3.00# rpm -qi zip-2.3-3
Name           : zip                      Relocations: /opt/freeware
Version        : 2.3                      Vendor: (none)
Release        : 3                        Build Date: Sat Nov 23 04:25:40 IST 2002
Install date:  Wed Dec 26 06:37:50 IST 2007      Build Host: emperor.aixplab.austin.ibm.com
Group          : Applications/Archiving      Source RPM: zip-2.3-3.src.rpm
Size           : 323779                    License: IBM_ILA
URL            : http://www.info-zip.org/pub/infozip/Zip.html
Summary        : A file compression and packaging utility compatible with PKZIP.
Description    :
The zip program is a compression and file packaging utility. Zip is
analogous to a combination of the UNIX tar and compress commands and
is compatible with PKZIP (a compression and file packaging utility for
MS-DOS systems).

Install the zip package if you need to compress files using the zip
program.
```

Install a rpm package

rpm -i invokes rpm in the install mode.

Typically **rpm -ivh** is used, v for verbose and h to display installation progress with hash symbol, like ftp transfer options. Sometimes **--ignoreos** flag is also useful. You can use the **--force** for force installation and **--nodeps** for avoid checking the dependencies (notice the double hyphen --).

```
bash-3.00# ls -l | grep zip*.rpm
-rw-r----- 1 root system 124474 Dec 04 22:18 zip-2.3-3.aix4.3.ppc.rpm
bash-3.00# rpm -ivh zip-2.3-3.aix4.3.ppc.rpm
zip
#####
bash-3.00#
```

```
bash-3.00# ls -l | grep zip*.rpm
-rw-r----- 1 root system 124474 Dec 04 22:18 zip-2.3-3.aix4.3.ppc.rpm
bash-3.00# rpm -ivh --ignoreos zip-2.3-3.aix4.3.ppc.rpm
zip
#####
bash-3.00#
```

Uninstall a rpm package

rpm -e invokes rpm in the erase mode.

```
bash-3.00# rpm -qa | grep zip
zip-2.3-3
bash-3.00# rpm -qa
cdrecord-1.9-7
mkisofs-1.13-4
zip-2.3-3
bash-3.0-1
bash-doc-3.0-1
AIX-rpm-5.3.0.60-2
vnc-3.3.3r2-6
bash-3.00# rpm -e zip
bash-3.00# rpm -qa
cdrecord-1.9-7
mkisofs-1.13-4
bash-3.0-1
bash-doc-3.0-1
AIX-rpm-5.3.0.60-2
vnc-3.3.3r2-6
bash-3.00#
```

Upgrade a rpm package.

rpm -U invokes rpm in the upgrade mode.

This mode can be used to upgrade an already installed package, else it installs the package, just like install mode.

```

bash-3.00# rpm -Uvh zip-2.3-3.aix4.3.ppc.rpm
package zip-2.3-3 is already installed
bash-3.00# rpm -e zip
bash-3.00# rpm -Uvh zip-2.3-3.aix4.3.ppc.rpm
zip #####
bash-3.00# rpm -qa
cdrecord-1.9-7
mkisofs-1.13-4
zip-2.3-3
bash-3.0-1
bash-doc-3.0-1
AIX-rpm-5.3.0.60-2
vnc-3.3.3r2-6
bash-3.00# █

```

Freshen a rpm package.

rpm -F

This will upgrade packages, but only if an earlier version currently exists.

```

bash-3.00# rpm -e zip
bash-3.00# rpm -Fvh zip-2.3-3.aix4.3.ppc.rpm
bash-3.00# rpm -qa
cdrecord-1.9-7
mkisofs-1.13-4
bash-3.0-1
bash-doc-3.0-1
AIX-rpm-5.3.0.60-2
vnc-3.3.3r2-6
bash-3.00# █

```

In AIX many of the non lpp programs are shipped in rpm formats.

You can refer the website <http://www.rpm.org/> for more information.

Employing alt_disk_install utility

Installs an alternate disk with a mksysb install image or clones the currently running system to an alternate disk. This command is obsolete in AIX 5.3.

In AIX 5.3 the alt_disk_install command has been broken up into three commands:

alt_disk_copy,
alt_disk_mksysb, and
alt_rootvg_op.

No new functionality will be added to this command.

```
# lslpp -l bos.alt_disk_install.rte
Fileset              Level  State      Description
-----
Path: /usr/lib/objrepos
bos.alt_disk_install.rte  5.3.0.63  APPLIED    Alternate Disk Installation
Runtime

# lslpp -l bos.alt_disk_install.boot_images
Fileset              Level  State      Description
-----
Path: /usr/lib/objrepos
bos.alt_disk_install.boot_images
                    5.3.0.62  APPLIED    Alternate Disk Installation
Disk Boot Images
```

```
bash-3.00# which alt_disk_install
/usr/sbin/alt_disk_install
bash-3.00# file /usr/sbin/alt_disk_install
/usr/sbin/alt_disk_install: shell script - ksh (Korn shell)
```

At the end of the install, a volume group, altinst_rootvg, is left on the target disks in the varied off state as a place holder. If varied on, it shows as owning no logical volumes, but it does indeed contain logical volumes, but they have been removed from the ODM because their names now conflict with the names of the logical volumes on the running system. It is recommended that you not vary on the altinst_rootvg volume group, but just leave the definition there as a place holder

After rebooting from the new alternate disk, the former rootvg volume group shows up in a lspv listing as "old_rootvg", and includes all disk(s) in the original rootvg. This former rootvg volume group is set to NOT varyon at reboot, and should ONLY be removed with the -X flag (i.e. alt_disk_install -X old_rootvg).

The AIX 4.3.1 and greater version of `alt_disk_install` can be executed in phases. The install is divided into three phases, and the default is to perform all three phases.

Phase 1

Creates the `altinst_rootvg` volume group, the `alt_` "logical volumes", the `/alt_inst` file systems, and restores the `mksysb` or `rootvg` data.

Phase 2

Runs any specified customization script, installs updates, new filesets, fixes or bundles (cloning only), copies a `resolv.conf` file if specified, and copies files over to remain a NIM client if specified.

Phase 3

Unmounts the `/alt_inst` file systems, renames the file systems and logical volumes, removes the `alt_` logical volumes, names ODM and varies off the `altinst_rootvg`. It sets the bootlist and reboots if specified.

You can run each phase separately, run Phases 1 and 2 together, or run Phases 2 and 3 together. Phase 2 can be run multiple times before Phase 3 is run.

You must run Phase 3 to get a volume group that is a usable `rootvg`. Running Phase 1 and 2 leave the `/alt_inst` file systems mounted.

If you have run Phase 1 and or Phase 2, and want to start over (remove the `altinst_rootvg`), run the `alt_disk_install-x` command to clean up.

Creating alt_disk_install configuration,

alt_disk_install -C -i /tmp/image.data.modified -V -O hdisk2

-V configuration can be omitted in case verbose messages are not required.

Remember to use -O flag if the target disk or disks will become the rootvg of a different system (such as in the case of logical partitioning or system disk swap). Performs a device reset on the target altinst_rootvg. This will cause alt_disk_install to NOT retain any user defined device configurations.

```

bash-3.00# ls -l /tmp/image.data.modified
-rw-r--r--  1 root      system      7948 Apr  9 15:07 /tmp/image.data.modified
bash-3.00# alt_disk_install -C -i /tmp/image.data.modified -O hdisk2
+-----+
ATTENTION: calling new module /usr/sbin/alt_disk_copy. Please see the
alt_disk_copy man page and documentation for more details.
Executing command: /usr/sbin/alt_disk_copy -i "/tmp/image.data.modified" -O -d "hdisk2"
+-----+
Checking disk sizes.
Creating cloned rootvg volume group and associated logical volumes.
Creating logical volume alt_hd5.
Creating logical volume alt_hd6.
Creating logical volume alt_hd8.
Creating logical volume alt_hd4.
Creating logical volume alt_hd2.
Creating logical volume alt_hd9var.
Creating logical volume alt_hd3.
Creating logical volume alt_hd1.
Creating logical volume alt_hd10opt.
Creating /alt_inst/ file system.
Creating /alt_inst/home file system.
Creating /alt_inst/opt file system.
Creating /alt_inst/tmp file system.
Creating /alt_inst/usr file system.
Creating /alt_inst/var file system.
Generating a list of files
for backup and restore into the alternate file system...
Backing-up the rootvg files and restoring them to the alternate file system...

```

The ***alt_disk_install -C*** invokes ***alt_disk_copy*** command in turn as shown below.

The data copy is getting done using the ***backup*** and ***restore*** commands not the ***cp*** command.

```

root 356386 544838  0 15:10:42 pts/0  0:00 backbyname -i -U -q -f -
root 364626 430286  1 15:10:42 pts/0  0:00 restbyname -q -f -
root 368698 544838  0 15:10:42 pts/0  0:00 cat /tmp/.include.list.196628
root 430286 196628  0 15:10:42 pts/0  0:00 /usr/bin/ksh /usr/sbin/alt_disk_copy -i /tmp/image.data.modifi
ed -O -d hdisk2
root 528606 323722  0 15:12:21 pts/1  0:00 ps -ef
root 544838 196628  0 15:10:42 pts/0  0:00 /usr/bin/ksh /usr/sbin/alt_disk_copy -i /tmp/image.data.modifi
ed -O -d hdisk2

```

As you can see below, the new disk is assigned with altinst_rootvg Volume group and the filesystems are all mounted with the /alt_inst prefix.

```
bash-3.00# lspv
hdisk0          00c28f209d5c7347          rootvg          active
hdisk1          00c28f20ad8044a6          rootvg          active
hdisk2          00c28f20328e9a37          altinst_rootvg active
bash-3.00# df -g
Filesystem      GB blocks      Free %Used      Iused %Iused Mounted on
/dev/hd4        1.00          0.97   4%       2772    2% /
/dev/hd2        4.00          2.12  47%      40760   8% /usr
/dev/hd9var     4.00          3.98   1%        443    1% /var
/dev/hd3        4.00          3.99   1%         95    1% /tmp
/dev/hd1        1.00          1.00   1%         5     1% /home
/proc           -             -      -         -     - /proc
/dev/hd10opt    1.00          0.92   8%      2249    2% /opt
/dev/alt_hd4    1.00          0.97   4%      2746    2% /alt_inst
/dev/alt_hd1    1.00          1.00   1%         5     1% /alt_inst/home
/dev/alt_hd10opt 1.00          0.92   8%      2249    2% /alt_inst/opt
/dev/alt_hd3    4.00          3.99   1%         88    1% /alt_inst/tmp
/dev/alt_hd2    4.00          3.81   5%      4363    1% /alt_inst/usr
/dev/alt_hd9var 4.00          4.00   1%         4     1% /alt_inst/var
```

```
bash-3.00# lsvg -l altinst_rootvg
altinst_rootvg:
LV NAME          TYPE      LPs   PPs   PVs  LV STATE      MOUNT POINT
alt_hd5          boot      1     1     1    closed/syncd  N/A
alt_hd6          paging    16    16    1    closed/syncd  N/A
alt_hd8          jfs2log   1     1     1    open/syncd    N/A
alt_hd4          jfs2      4     4     1    open/syncd    /alt_inst
alt_hd2          jfs2     16    16    1    open/syncd    /alt_inst/usr
alt_hd9var       jfs2     16    16    1    open/syncd    /alt_inst/var
alt_hd3          jfs2     16    16    1    open/syncd    /alt_inst/tmp
alt_hd1          jfs2      4     4     1    open/syncd    /alt_inst/home
alt_hd10opt      jfs2      4     4     1    open/syncd    /alt_inst/opt
```

```

bash-3.00# df -g
Filesystem      GB blocks      Free %Used      Iused %Iused Mounted on
/dev/hd4         1.00          0.97   4%       2772    2% /
/dev/hd2         4.00          2.12  47%      40760   8% /usr
/dev/hd9var      4.00          3.98   1%        443    1% /var
/dev/hd3         4.00          3.99   1%         95    1% /tmp
/dev/hd1         1.00          1.00   1%         5     1% /home
/proc            -             -      -         -     - /proc
/dev/hd10opt     1.00          0.92   8%       2249    2% /opt
/dev/alt_hd4     1.00          0.97   4%       2746    2% /alt_inst
/dev/alt_hd1     1.00          1.00   1%         5     1% /alt_inst/home
/dev/alt_hd10opt 1.00          0.92   8%       2249    2% /alt_inst/opt
/dev/alt_hd3     4.00          3.99   1%         88    1% /alt_inst/tmp
/dev/alt_hd2     4.00          2.69  33%      30177   5% /alt_inst/usr
/dev/alt_hd9var  4.00          4.00   1%         4     1% /alt_inst/var

```

Finishing the alt_disk_install command

```

Generating a list of files
for backup and restore into the alternate file system...
Backing-up the rootvg files and restoring them to the alternate file system...
Modifying ODM on cloned disk.
Building boot image on cloned disk.
Resetting all device attributes.
NOTE: The first boot from altinst_rootvg will prompt to define the new
system console.
Resetting all device attributes.
NOTE: The first boot from altinst_rootvg will prompt to define the new
system console.
forced unmount of /alt_inst/var
forced unmount of /alt_inst/usr
forced unmount of /alt_inst/tmp
forced unmount of /alt_inst/opt
forced unmount of /alt_inst/home
forced unmount of /alt_inst
forced unmount of /alt_inst
Changing logical volume names in volume group descriptor area.
Fixing IV control blocks...
Fixing file system superblocks...
Bootlist is set to the boot disk: hdisk2

```

After successful creation of the alt_inst_disk, remove the hdisk from the production system using the rmdev command and also physically disconnect the hdisk. After this in the production system clean up the alt_disk_install volume group configuration.

Clean up the alt_disk_install volume group configurations from the production system.

```
bash-3.00# lsvg
rootvg
altinst_rootvg
bash-3.00# alt_disk_install -X altinst_rootvg
+-----+
ATTENTION: calling new module /usr/sbin/alt_rootvg_op. Please see the
alt_rootvg_op man page and documentation for more details.
Executing command: /usr/sbin/alt_rootvg_op -X altinst_rootvg
+-----+
Bootlist is set to the boot disk: hdisk1
bash-3.00# lsvg
rootvg
```

Target system after booting up from the alt_disk_install target disk.

```
# lspv
hdisk0          00c28f20328e9a37          rootvg          active
# lsvg -l rootvg
rootvg:
LV NAME          TYPE          LPs    PPs    PVs    LV STATE      MOUNT POINT
hd5              boot          1      1      1      closed/syncd  N/A
hd6              paging        16     16     1      open/syncd    N/A
hd8              jfs2log      1      1      1      open/syncd    N/A
hd4              jfs2         4      4      1      open/syncd    /
hd2              jfs2         16     16     1      open/syncd    /usr
hd9var           jfs2         16     16     1      open/syncd    /var
hd3              jfs2         16     16     1      open/syncd    /tmp
hd1              jfs2         4      4      1      open/syncd    /home
hd10opt          jfs2         4      4      1      open/syncd    /opt
#
```

Note that we got a clean system, with rootvg on hdisk0 with a pvid the source server assigned, but the device name changed from hdisk2 to hdisk0. Moreover the source is with mirrored logical volumes but the target is without mirroring.

Only left out configuration is just to set the bootlist to hdisk0.

Very Important : Since the flag -O clears all the custom device configuration, please do remember to put back the custom device configurations back after the new AIX bootup is completed. Most of them are for sys0, aio0, fscsiX, enX etc devices. But schedo, vmo, ioo, no, nsfo tuning parameters remain unchanged.

Modifying image.data file is very helpful at the time of restoration using nim. Most of the flexibility pertains to the fact that we can shrink filesystems, delete paging spaces and reduce the number of copies of logical volumes.

Filesystem Quota Management.

The disk quota system allows system administrators to control the number of files and data blocks that can be allocated to users or groups.

Disk quota system concept

The disk quota system, based on the Berkeley Disk Quota System, provides an effective way to control the use of disk space. The quota system can be defined for individual users or groups, and is maintained for each journaled file system (JFS and JFS2). The disk quota system establishes limits based on the following parameters that can be changed with the **edquota** command for JFS file systems and the **j2edlimit** command for JFS2 file systems:

- User's or group's soft limits
- User's or group's hard limits
- Quota grace period

The *soft limit* defines the number of 1 KB disk blocks or files under which the user must remain. The *hard limit* defines the maximum amount of disk blocks or files the user can accumulate under the established disk quotas. The *quota grace period* allows the user to exceed the soft limit for a short period of time (the default value is one week). If the user fails to reduce usage below the soft limit during the specified time, the system will interpret the soft limit as the maximum allocation allowed, and no further storage is allocated to the user. The user can reset this condition by removing enough files to reduce usage below the soft limit.

```

$ lslpp -w /usr/sbin/quota
File                               Fileset                             Type
-----
/usr/sbin/quota                    bos.sysmgt.quota                    File
$ lslpp -f bos.sysmgt.quota
Fileset                             File
-----
Path: /usr/lib/objrepos
bos.sysmgt.quota 5.3.0.60
                               /usr/sbin/repquota
                               /usr/sbin/edquota
                               /usr/sbin/quot
                               /usr/sbin/quotacheck
                               /usr/sbin/quotaoff -> /usr/sbin/quotaon
                               /usr/sbin/quotaon
                               /usr/sbin/j2edlimit
                               /usr/sbin/quota

```

Determine which file systems require quotas. The disk quota system can be used only with the journaled file system. Do not establish disk quotas for the **/tmp** file system.

Note: Because many editors and system utilities create temporary files in the **/tmp** file system, it must be free of quotas.

Use the **chfs** command to include the **userquota** and **groupquota** quota configuration attributes in the **/etc/filesystems** file. The following example uses the **chfs** command to enable user quotas on the **/ora** file system:

```
chfs -a "quota = userquota" /ora
```

To enable both user and group quotas on the **/home** file system, type:

```
chfs -a "quota = userquota,groupquota" /ora
```

Alternately Use smitty chfs

```

[Entry Fields]
File system name           /ora
NEW mount point           [/ora]
SIZE of file system
    Unit Size              512bytes
    Number of units        [2097152]
Mount GROUP               []
Mount AUTOMATICALLY at system restart?  yes
PERMISSIONS               read/write
Mount OPTIONS             []
Start Disk Accounting?    no
Block Size (bytes)        4096
Inline Log?                no
Inline Log size (MBytes)  [0]
Extended Attribute Format  [v1]
ENABLE Quota Management?  all
Allow Small Inode Extents? no

```

```

Command: OK          stdout: yes          stderr: no

Before command completion, additional instructions may appear below.

Initializing quota file /ora/quota.user
Initializing quota file /ora/quota.group
*** Checking user and group quotas for /dev/fslv00 (/ora)
root      fixed: inodes 1 -> 4   blocks 32 -> 64
system    fixed: inodes 0 -> 4   blocks 0 -> 64

```

```

bash-3.00# cd /ora
bash-3.00# pwd
/ora
bash-3.00# ls -la
total 136
drwxr-xr-x   3 root    system      256 May 10 18:36 .
drwxr-xr-x  37 root    system     1536 May 08 20:56 ..
drwxr-xr-x   2 root    system      256 Mar 26 16:54 lost+found
-rw-r-----   1 root    system    32768 May 10 18:36 quota.group
-rw-r-----   1 root    system    32768 May 10 18:36 quota.user
bash-3.00# file quota*
quota.group: commands text
quota.user:  commands text
bash-3.00# █

```

```

bash-3.00# grep -p ora /etc/filesystems
/ora:
    dev          = /dev/fslv00
    vfs          = jfs2
    log          = /dev/loglv02
    mount        = true
    options      = rw
    account      = false
    quota        = userquota,groupquota

```

```

/dev/fslv00  --      /ora          jfs2 2097152 rw      yes no
(lv size: 2097152, fs size: 2097152, block size: 4096, sparse files: yes, inline log: no, inline log size: 0, EAformat: v1,
Quota: userquota,groupquota, DMAPi: no, VIX: no)

```

Optionally, specify alternate disk quota file names. The **quota.user** and **quota.group** file names are the default names located at the root directories of the file systems enabled with quotas.

```
chfs -a "userquota = /ora/myquota.user" -a "groupquota = /ora/myquota.group" /ora
```

Use the *edquota* command to create and edit quotas for JFS file systems.
To manage quotas on a JFS2 file system, use the *j2edlimit* command.

As most of the filesystems nowadays is jfs2 our further discussions are limited to the *j2edlimit* command.

Quotas are managed in JFS2 file systems through the use of **Limits Classes**. Each Limits Class has hard and soft limits for disk space and file, and grace periods for exceeding the soft limits. Individual users and groups may be assigned to a Limits Class and are then subject to the quotas defined by that class. Any user or group not assigned to a class is subject to the quotas defined by the **default class** (Class ID 0).

Fields required for quota definition are the following :

Block Hard Limit : The total amount of 1KB blocks the user or group will be allowed to use, including temporary storage during a quota grace period.

Block Soft Limit : The number of 1KB blocks the user or group will be allowed to use during normal operations.

File Hard Limit : The total number of files the user or group will be allowed to create, including temporary files created during a quota grace period.

File Soft Limit : The number of files the user or group will be allowed to create during normal operations.

Block Grace Period : Amount of time a user can exceed the Block Soft Limit before it becomes enforced as a hard limit.

File Grace Period : Amount of time a user can exceed the File Soft Limit before it becomes enforced as a hard limit.

```
bash-3.00# j2edlimit
Usage:  j2edlimit [-e] [-u|-g] Filesystem
        j2edlimit -l [-u|-g] Filesystem
        j2edlimit -d limID [-u|-g] Filesystem
        j2edlimit -a limID {-u UID | -g GID} Filesystem
        j2edlimit -c limID [-u|-g] Filesystem
```

To list the current user and group quota allocation, use the following commands :

```
bash-3.00# j2edlimit -l -g /ora
Group Limits Classes for file system /ora
Block Limits units: g=gigabyte, m=megabyte, or optional k=kilobyte
Grace Period units: d=days, h=hours, m=minutes, or s=seconds
Limits Class ID 0 is the default class.

  ID      Block Limits          File Limits          Grace Period
  soft    hard      soft    hard      block    file
  0        0         0        0         0         0
bash-3.00# j2edlimit -l -u /ora
User Limits Classes for file system /ora
Block Limits units: g=gigabyte, m=megabyte, or optional k=kilobyte
Grace Period units: d=days, h=hours, m=minutes, or s=seconds
Limits Class ID 0 is the default class.

  ID      Block Limits          File Limits          Grace Period
  soft    hard      soft    hard      block    file
  0        0         0        0         0         0
```

Use `j2edlimit -e <filesystem>` to edit the Limits Classes for the file system specified on the command line.

```
bash-3.00# j2edlimit -e -u /ora
"/tmp/J2EdLim.JMlfqa" 11 lines, 509 characters
User Limits Classes for file system /ora
Block Limits units: g=gigabyte, m=megabyte, or optional k=kilobyte
Grace Period units: d=days, h=hours, m=minutes, or s=seconds
Limits Class ID 0 is the default class.
Prepend '-' to ID to delete Limits Class.
Use '+' for ID to add new Limits Class.
```

ID	Block Limits		File Limits		Grace Period	
	soft	hard	soft	hard	block	file
0	10m	20m	10	20	1d	1d

Defining another Limit class :

```
bash-3.00# j2edlimit -e -u /ora
"/tmp/J2EdLim.7Ur-ia" 11 lines, 509 characters
User Limits Classes for file system /ora
Block Limits units: g=gigabyte, m=megabyte, or optional k=kilobyte
Grace Period units: d=days, h=hours, m=minutes, or s=seconds
Limits Class ID 0 is the default class.
Prepend '-' to ID to delete Limits Class.
Use '+' for ID to add new Limits Class.
```

ID	Block Limits		File Limits		Grace Period	
	soft	hard	soft	hard	block	file
0	10m	20m	10	20	1d	2d
+	5m	7m	5	7	1h	2h

```
bash-3.00# j2edlimit -l -u /ora
User Limits Classes for file system /ora
Block Limits units: g=gigabyte, m=megabyte, or optional k=kilobyte
Grace Period units: d=days, h=hours, m=minutes, or s=seconds
Limits Class ID 0 is the default class.
```

ID	Block Limits		File Limits		Grace Period	
	soft	hard	soft	hard	block	file
0	10m	20m	10	20	1d	2d
1	5m	7m	5	7	1h	2h

```
bash-3.00#
```

To change the default Limits Class :

Please remember that making any Class as default is overwriting the Limits Class 0 with the selected Class.

```
bash-3.00# j2edlimit -l -u /ora
User Limits Classes for file system /ora
Block Limits units: g=gigabyte, m=megabyte, or optional k=kilobyte
Grace Period units: d=days, h=hours, m=minutes, or s=seconds
Limits Class ID 0 is the default class.
```

ID	Block Limits		File Limits		Grace Period	
	soft	hard	soft	hard	block	file
0	5m	7m	5	7	1h	2h
1	5m	7m	5	7	1h	2h
2	10m	20m	10	20	1d	2d

```
bash-3.00# j2edlimit -d 2 /ora
User Limits Classes for file system /ora
Block Limits units: g=gigabyte, m=megabyte, or optional k=kilobyte
Grace Period units: d=days, h=hours, m=minutes, or s=seconds
Limits Class ID 0 is the default class.
```

ID	Block Limits		File Limits		Grace Period	
	soft	hard	soft	hard	block	file
0	10m	20m	10	20	1d	2d

```
bash-3.00# j2edlimit -l -u /ora
User Limits Classes for file system /ora
Block Limits units: g=gigabyte, m=megabyte, or optional k=kilobyte
Grace Period units: d=days, h=hours, m=minutes, or s=seconds
Limits Class ID 0 is the default class.
```

ID	Block Limits		File Limits		Grace Period	
	soft	hard	soft	hard	block	file
0	10m	20m	10	20	1d	2d
1	5m	7m	5	7	1h	2h
2	10m	20m	10	20	1d	2d

```
bash-3.00#
```

Assigning the associated User to the specified Limits Class.

```

bash-3.00# j2edlimit -l -u /ora
User Limits Classes for file system /ora
Block Limits units: g=gigabyte, m=megabyte, or optional k=kilobyte
Grace Period units: d=days, h=hours, m=minutes, or s=seconds
Limits Class ID 0 is the default class.

  ID      Block Limits          File Limits          Grace Period
  soft    hard      soft    hard      block   file
0       10m     20m     10      20       1d      2d
1        5m     7m      5       7       1h      2h
2       10m     20m     10      20       1d      2d
bash-3.00# j2edlimit -a 1 -u dinil /ora
User Limits Classes for file system /ora
Block Limits units: g=gigabyte, m=megabyte, or optional k=kilobyte
Grace Period units: d=days, h=hours, m=minutes, or s=seconds

  ID      Block Limits          File Limits          Grace Period
  soft    hard      soft    hard      block   file
1        5m     7m      5       7       1h      2h
bash-3.00# j2edlimit -l -u /ora
User Limits Classes for file system /ora
Block Limits units: g=gigabyte, m=megabyte, or optional k=kilobyte
Grace Period units: d=days, h=hours, m=minutes, or s=seconds
Limits Class ID 0 is the default class.

  ID      Block Limits          File Limits          Grace Period
  soft    hard      soft    hard      block   file
0       10m     20m     10      20       1d      2d
1        5m     7m      5       7       1h      2h
2       10m     20m     10      20       1d      2d

```

Enable the quota system with the **quotaon** command. The **quotaon** command enables quotas for a specified file system, or for all file systems with quotas (as indicated in the `/etc/filesystems` file) when used with the `-a` flag.

```

bash-3.00# quotaon
Usage:
    quotaon [-g] [-u] [-v] -a
    quotaon [-g] [-u] [-v] filesystem
bash-3.00# quotaon /ora
bash-3.00# █

```

Checking the quota implementation

```

$ whoami
dinil
$ cd /ora
$ touch test1
touch: 0652-046 Cannot create test1.
$ su -
root's Password:
3004-501 Cannot su to "root" : Authentication is denied.
$ su -
root's Password:
# bash
bash-3.00# chmod -R 777 /ora
bash-3.00# chmod -R 777 /ora
bash-3.00# exit
#
$ touch test1
$ touch test2
$ touch test3
$ touch test4
$ touch test5
$ touch test6

/ora: warning, user inode quota exceeded
$ touch test7
$ touch test8
touch: 0652-046 Cannot create test8.

/ora: operation failed, user inode quota limit reached

```

```

$ ls -l
total 136
-rwxrwxrwx  1 root    system      414 May 10 22:01 HISTFILE.LOG
drwxrwxrwx  2 root    system      256 Mar 26 16:54 lost+found
-rwxrwxrwx  1 root    system    32768 May 12 00:57 quota.group
-rwxrwxrwx  1 root    system    32768 May 12 00:57 quota.user
-rw-r-----  1 dinil   staff         0 May 12 00:59 test1
-rw-r-----  1 dinil   staff         0 May 12 00:59 test2
-rw-r-----  1 dinil   staff         0 May 12 00:59 test3
-rw-r-----  1 dinil   staff         0 May 12 01:00 test4
-rw-r-----  1 dinil   staff         0 May 12 01:00 test5
-rw-r-----  1 dinil   staff         0 May 12 01:00 test6
-rw-r-----  1 dinil   staff         0 May 12 01:00 test7

```

The **quota** command displays disk usage and quotas.

```
quota [ -u [ User ] ] [ -g [ Group ] ] [ -v | -q ]
```

```
bash-3.00# quota -u dinil
Disk quotas for user dinil (uid 207):
  Filesystem  blocks    quota   limit   grace   files   quota   limit   grace
   /ora        0      5120   7168         6 *     5       7     1:54
bash-3.00# quota -v
Disk quotas for user root (uid 0):
  Filesystem  blocks    quota   limit   grace   files   quota   limit   grace
   /ora       72        0       0         5       0       0
```

```
bash-3.00# quotacheck -v /ora
*** Checking user and group quotas for /dev/fslv00 (/ora)
dinil    fixed: inodes 6 -> 7
staff    fixed: inodes 6 -> 7
```

The **repquota** command prints a summary of quotas and disk usage for a file system specified by the FileSystem parameter.

```
repquota [ -v ] [ -c ] [ -g ] [ -u ] [ -l ] { -a | FileSystem ... }
```

```
bash-3.00# repquota -av
*** Report for group quotas on /ora (/dev/fslv00)
      Block limits
Group  --  used  soft  hard  grace  used  soft  hard  grace
system --   72    0    0         5     0    0
staff  --    0    0    0         7     0    0

*** Report for user quotas on /ora (/dev/fslv00)
      Block limits
User   --  used  soft  hard  grace  used  soft  hard  grace
root   --   72    0    0         5     0    0
dinil  -+    0  5120  7168         7     5     7     1:49
bash-3.00# █
```

Multiple Instances of AIX on a Single Root Volume Group

The multibos command allows the root level administrator to create multiple instances of AIX on the same rootvg. The multibos setup operation creates a standby Base Operating System (BOS) that boots from a distinct boot logical volume (BLV). This creates two bootable sets of BOS on a given rootvg. The administrator can boot from either instance of BOS by specifying the respective BLV as an argument to the bootlist command or using system firmware boot operations. Two bootable instances of BOS can be simultaneously maintained. The instance of BOS associated with the booted BLV is referred to as the active BOS. The instance of BOS associated with the BLV that has not been booted is referred to as the standby BOS. Currently, only two instances of BOS are supported per rootvg.

The multibos command allows the administrator to access, install maintenance and technology levels for, update, and customize the standby BOS either during setup or in subsequent customization operations. Installing maintenance and technology updates to the standby BOS does not change system files on the active BOS. This allows for concurrent update of the standby BOS, while the active BOS remains in production.

In addition, the multibos command has the ability to copy or share logical volumes and file systems. By default, the BOS file systems (currently /, /usr, /var, /opt, and /home) and the boot logical volume are copied. The administrator can make copies of additional BOS objects (using the -L flag).

All other file systems and logical volumes are shared between instances of BOS. Separate log device logical volumes (for example, those that are not contained within the file system) are not supported for copy and will be shared.

multibos : - command that creates, updates, and manages multiple versions of the BOS on a rootvg. "Only the root user can run the multibos command."

Syntax

multibos -s [-l Device {-a | -f File | -b File}] [-e File] [-i File] [-L File] [-pntNX] - Creates an instance of standby BOS.

multibos -c [-l Device {-a | -f File | -b File}] [-pnNX] - Performs a customized update of the software in standby BOS.

multibos -m [-pnX] - Mounts standby BOS.

multibos -u [-pnX] - Unmounts standby BOS.

multibos -B [-ntX] - Build boot image operation. The standby boot image is created and written to the standby BLV using the AIX bosboot command.

multibos -S [-nX] - Initiates an interactive shell with chroot access to the standby BOS file systems.

multibos -R [-ptX] - Removes all standby BOS objects.

-n Does not perform cleanup upon failure.

-p Performs a preview of the given operation.

-X Allows for automatic file system expansion if space is needed to perform tasks related to multibos.

- The multibos command is supported on AIX 5L Version 5.3 with the 5300-03 ML and later.
- The current rootvg must have enough space for each BOS object copy.
- BOS object copies are placed on the same disk or disks as the original.
- The total number of copied logical volumes cannot exceed 128.
- The total number of copied logical volumes and shared logical volumes are subject to VG limits.

Creates an instance of standby BOS

```

bash-3.00# multibos -Xs
Initializing multibos methods ...
Initializing log /etc/multibos/logs/op.alog ...
Gathering system information ...

+-----+
Setup Operation
+-----+
Verifying operation parameters ...
Creating image.data file ...

+-----+
Logical Volumes
+-----+
Creating standby BOS logical volume bos_hd5
Creating standby BOS logical volume bos_hd4
Creating standby BOS logical volume bos_hd2
Creating standby BOS logical volume bos_hd9var
Creating standby BOS logical volume bos_hd10opt

+-----+
File Systems
+-----+
Creating all standby BOS file systems ...
Creating standby BOS file system /bos_inst on logical volume bos_hd4
Creating standby BOS file system /bos_inst/usr on logical volume bos_hd2
Creating standby BOS file system /bos_inst/var on logical volume bos_hd9var
Creating standby BOS file system /bos_inst/opt on logical volume bos_hd10opt

+-----+
Mount Processing
+-----+
Mounting all standby BOS file systems ...
Mounting /bos_inst
Mounting /bos_inst/usr
Mounting /bos_inst/var
Mounting /bos_inst/opt

+-----+
BOS Files
+-----+
Including files for file system /
Including files for file system /usr
Including files for file system /var
Including files for file system /opt

```

```

Copying files using backup/restore utilities ...
Percentage of files copied: 0.00%
Percentage of files copied: 1.66%
Percentage of files copied: 3.33%
Percentage of files copied: 5.00%
Percentage of files copied: 6.67%
Percentage of files copied: 8.34%
Percentage of files copied: 10.00%
Percentage of files copied: 11.67%
Percentage of files copied: 13.34%
Percentage of files copied: 15.01%
Percentage of files copied: 16.68%
Percentage of files copied: 18.35%
Percentage of files copied: 20.01%
Percentage of files copied: 21.68%
Percentage of files copied: 23.35%
Percentage of files copied: 25.02%
Percentage of files copied: 26.69%
Percentage of files copied: 28.36%
Percentage of files copied: 30.02%
Percentage of files copied: 31.69%
Percentage of files copied: 33.36%
Percentage of files copied: 35.03%
Percentage of files copied: 36.70%
Percentage of files copied: 38.37%
Percentage of files copied: 40.03%
Percentage of files copied: 41.70%
Percentage of files copied: 43.37%

```

Building the Standby BOS uses, backbyname and restbyname commands.

```

root 446604 532654 0 11:36:49 pts/0 0:00 bash
root 512090 446604 2 13:23:37 pts/0 0:01 /usr/bin/ksh /usr/sbin/multibos -Xs
root 516314 401436 0 13:24:59 pts/2 0:00 -ksh
root 520238 516314 0 13:25:13 pts/2 0:00 ps -ef
root 524344 536656 1 13:24:06 pts/0 0:00 restbyname -q -f -
root 528500 512090 0 13:24:06 pts/0 0:00 /usr/bin/ksh /usr/sbin/multibos -Xs
root 532654 438296 0 11:36:29 pts/0 0:00 -ksh
root 536656 512090 0 13:24:06 pts/0 0:00 /usr/bin/ksh /usr/sbin/multibos -Xs
root 540772 512090 0 13:23:38 pts/0 0:00 /usr/bin/alog -f /etc/multibos/logs/op.alog -s 8388608
root 553170 528500 0 13:24:06 pts/0 0:00 backbyname -i -q -f -

```

These new BOS areas are mounted under /bos_inst root.

```

# df -g
Filesystem      GB blocks      Free %Used      Iused %Iused Mounted on
/dev/hd4        1.00           0.96   4%         2779    2% /
/dev/hd2        4.00           2.12  47%        40760   8% /usr
/dev/hd9var     4.00           3.98   1%          449    1% /var
/dev/hd3        4.00           3.99   1%          110    1% /tmp
/dev/hd1        1.00           1.00   1%           5     1% /home
/proc           -              -      -           -     - /proc
/dev/hd10opt    1.00           0.92   8%         2249    2% /opt
/dev/swdump     15.00          3.15  80%         5308    1% /swdump
/dev/nim1       6.00           3.77  38%           5     1% /nim1
/dev/bos_hd4    1.00           0.96   4%         2757    2% /bos_inst
/dev/bos_hd2    4.00           3.98   1%          563    1% /bos_inst/usr
/dev/bos_hd9var 4.00           4.00   1%           4     1% /bos_inst/var
/dev/bos_hd10opt 1.00           0.92   8%         2249    2% /bos_inst/opt

```

Towards the end of the multibos command the bosboot will be run. Interesting syntax employed in the bosboot command.

```

root 512090 446604 0 13:23:37 pts/0 0:25 /usr/bin/ksh /usr/sbin/multibos -Xs
root 516314 401436 0 13:24:59 pts/2 0:00 -ksh
root 520224 516314 0 13:37:32 pts/2 0:00 bash
root 524418 536830 71 13:42:53 pts/0 0:00 /usr/sbin/mkram /tmp/bosboot_536830_8822/Bootram.fs
root 528634 520224 1 13:42:54 pts/2 0:00 ps -ef
root 532654 438296 0 11:36:29 pts/0 0:00 -ksh
root 536830 512090 0 13:42:44 pts/0 0:00 /usr/bin/ksh /usr/sbin/bosboot -ad hdisk0 -l bos_hd5 -M standby
root 540772 512090 0 13:23:38 pts/0 0:00 /usr/bin/alog -f /etc/multibos/logs/op.alog -s 8388608
root 553174 520224 0 13:42:54 pts/2 0:00 more

```

```

Percentage of files copied: 83.41%
Percentage of files copied: 85.08%
Percentage of files copied: 86.74%
Percentage of files copied: 88.41%
Percentage of files copied: 90.08%
Percentage of files copied: 91.75%
Percentage of files copied: 93.42%
Percentage of files copied: 95.09%
Percentage of files copied: 96.75%
Percentage of files copied: 98.42%
Percentage of files copied: 100.00%

```

```

+-----+
Boot Partition Processing
+-----+
Active boot logical volume is hd5.
Standby boot logical volume is bos_hd5.
Creating standby BOS boot image on boot logical volume bos_hd5
bosboot: Boot image is 33193 512 byte blocks.

+-----+
Mount Processing
+-----+
Unmounting all standby BOS file systems ...
Unmounting /bos_inst/opt
Unmounting /bos_inst/var
Unmounting /bos_inst/usr
Unmounting /bos_inst

+-----+
Bootlist Processing
+-----+
Verifying operation parameters ...
Setting bootlist to logical volume bos_hd5 on hdisk0.
ATTENTION: firmware recovery string for BLV ():
boot

Log file is /etc/multibos/logs/op.alog
Return Status = SUCCESS
bash-3.00# █

```

```
bash-3.00# bootlist -m normal -o
hdisk0 blv=bos_hd5
hdisk0 blv=hd5
```

```
bash-3.00# lsvg -l rootvg
rootvg:
LV NAME          TYPE          LPs    PPs    PVs    LV STATE      MOUNT POINT
hd5              boot          1      1      1      closed/syncd  N/A
hd6              paging        16     16     1      open/syncd    N/A
hd8              jfs2log       1      1      1      open/syncd    N/A
hd4              jfs2          4      4      1      open/syncd    /
hd2              jfs2          16     16     1      open/syncd    /usr
hd9var           jfs2          16     16     1      open/syncd    /var
hd3              jfs2          16     16     1      open/syncd    /tmp
hd1              jfs2          4      4      1      open/syncd    /home
hd10opt          jfs2          4      4      1      open/syncd    /opt
bos_hd5          boot          1      1      1      closed/syncd  N/A
bos_hd4          jfs2          4      4      1      closed/syncd  /bos_inst
bos_hd2          jfs2          16     16     1      closed/syncd  /bos_inst/usr
bos_hd9var       jfs2          16     16     1      closed/syncd  /bos_inst/var
bos_hd10opt      jfs2          4      4      1      closed/syncd  /bos_inst/opt
```

```
bash-3.00# pwd
/etc/multibos/logs
bash-3.00# ls -l
total 16392
-rw-----  1 root      system    8388608 May 17 15:39 op.alog
-rw-----  1 root      system    115 May 17 15:39 scriptlog.080517153858.txt.Z
```

Use the command `alog -o -f op.log` to read the file.

```
bash-3.00# alog -o -f op.alog | more
=====
DATE: 2008.05.17.13:23:38 ID: [3 5 00C28F204C00] COMMAND: (multibos -Xs)
=====
Gathering system information ...

+-----+
Setup Operation
+-----+
Verifying operation parameters ...
Creating image.data file ...
```

Mounting the Standby BOS.

```

bash-3.00# multibos -mX
Initializing multibos methods ...
Initializing log /etc/multibos/logs/op.alog ...
Gathering system information ...

+-----+
BOS Mount Operation
+-----+
Verifying operation parameters ...

+-----+
Mount Processing
+-----+
Mounting all standby BOS file systems ...
Mounting /bos_inst
Mounting /bos_inst/usr
Mounting /bos_inst/var
Mounting /bos_inst/opt

Log file is /etc/multibos/logs/op.alog
Return Status = SUCCESS

```

```

bash-3.00# df -g
Filesystem      GB blocks      Free %Used      Iused %Iused Mounted on
/dev/hd4         1.00          0.96   4%        2778    2% /
/dev/hd2         4.00          2.12  47%       40760   8% /usr
/dev/hd9var      4.00          3.98   1%         449    1% /var
/dev/hd3         4.00          3.99   1%          95    1% /tmp
/dev/hd1         1.00          1.00   1%           5    1% /home
/proc            -              -      -          -      - /proc
/dev/hd10opt     1.00          0.92   8%        2249    2% /opt
/dev/swdump     15.00          3.15  80%        5308    1% /swdump
/dev/nim1        6.00          3.77  38%          5    1% /nim1
/dev/bos_hd4     1.00          0.96   4%        2758    2% /bos_inst
/dev/bos_hd2     4.00          2.12  47%       40760   8% /bos_inst/usr
/dev/bos_hd9var  4.00          3.98   1%         438    1% /bos_inst/var
/dev/bos_hd10opt 1.00          0.92   8%        2249    2% /bos_inst/opt

```

Unmounting the standby BOS.

```

bash-3.00# multibos -uX
Initializing multibos methods ...
Initializing log /etc/multibos/logs/op.alog ...
Gathering system information ...

+-----+
BOS Unmount Operation
+-----+
Verifying operation parameters ...

+-----+
Mount Processing
+-----+
Unmounting all standby BOS file systems ...
Unmounting /bos_inst/opt
Unmounting /bos_inst/var
umount: error unmounting /dev/bos_hd9var: Device busy
multibos: 0645-007 ATTENTION: unmount_dev() returned an unexpected result.
multibos: 0565-028 Error removing file systems.

Log file is /etc/multibos/logs/op.alog
Return Status: FAILURE
bash-3.00# pwd
/bos_inst/var
bash-3.00# cd
bash-3.00# █

```

```

bash-3.00# multibos -uX
Initializing multibos methods ...
Initializing log /etc/multibos/logs/op.alog ...
Gathering system information ...

+-----+
BOS Unmount Operation
+-----+
Verifying operation parameters ...

+-----+
Mount Processing
+-----+
Unmounting all standby BOS file systems ...
Unmounting /bos_inst/var
Unmounting /bos_inst/usr
Unmounting /bos_inst

Log file is /etc/multibos/logs/op.alog
Return Status = SUCCESS
█

```

Interacting with the standby BOS.

```

bash-3.00# multibos -SX
Initializing multibos methods ...
Initializing log /etc/multibos/logs/op.alog ...
Gathering system information ...

+-----+
Multibos Shell Operation
+-----+
Verifying operation parameters ...

+-----+
Mount Processing
+-----+
Mounting all standby BOS file systems ...
Mounting /bos_inst
Mounting /bos_inst/usr
Mounting /bos_inst/var
Mounting /bos_inst/opt

+-----+
Multibos Root Shell
+-----+
Starting multibos root shell ...
Active boot logical volume is hd5.
Script started, file is /etc/multibos/logs/scriptlog.080517135722.txt
MULTIBOS> █

```

```

MULTIBOS> df -g
Filesystem      GB blocks      Free %Used      Iused %Iused Mounted on
/dev/hd4        1.00          0.96   4%       2765    2% /
/dev/hd2        4.00          2.12  47%      40765   8% /usr
/dev/hd9var     4.00          3.98   1%        438    1% /var
/dev/hd3        4.00          3.99   1%        101    1% /tmp
/dev/hd1        1.00          0.96   4%       2765    2% /home
/proc           1.00          0.96   4%       2765    2% /proc
/dev/hd10opt    1.00          0.92   8%       2249    2% /opt
/dev/swdump     1.00          0.96   4%       2765    2% /swdump
/dev/nim1       1.00          0.96   4%       2765    2% /nim1
/dev/bos_hd4    1.00          0.96   4%       2765    2% /bos_inst
/dev/bos_hd2    4.00          2.12  47%      40765   8% /bos_inst/usr
/dev/bos_hd9var  4.00          3.98   1%        438    1% /bos_inst/var
/dev/bos_hd10opt 1.00          0.92   8%       2249    2% /bos_inst/opt
/usr/lib        4.00          2.12  47%      40760   8% /bos_inst/usr/lib/multibos_chroot/usr/lib
/usr/ccs/lib    4.00          2.12  47%      40760   8% /bos_inst/usr/lib/multibos_chroot/usr/ccs/lib
/tmp           4.00          3.99   1%        101    1% /bos_inst/tmp
MULTIBOS> hostname
activebos

```

```
MULTIBOS> mount
```

node	mounted	mounted over	vfs	date	options
/dev/hd4	/		jfs2	Apr 09 15:58	rw,log=/dev/hd8
/dev/hd2	/usr		jfs2	Apr 09 15:58	rw,log=/dev/hd8
/dev/hd9var	/var		jfs2	Apr 09 15:58	rw,log=/dev/hd8
/dev/hd3	/tmp		jfs2	Apr 09 15:58	rw,log=/dev/hd8
/dev/hd1	/home		jfs2	Apr 09 15:59	rw,log=/dev/hd8
/proc	/proc		procfs	Apr 09 15:59	rw
/dev/hd10opt	/opt		jfs2	Apr 09 15:59	rw,log=/dev/hd8
/dev/swdump	/swdump		jfs2	Apr 09 18:38	rw,log=/dev/loglv00
/dev/nim1	/nim1		jfs2	Apr 09 19:18	rw,log=/dev/loglv00
/dev/bos_hd4	/bos_inst		jfs2	May 17 13:57	rw,log=/dev/hd8
/dev/bos_hd2	/bos_inst/usr		jfs2	May 17 13:57	rw,log=/dev/hd8
/dev/bos_hd9var	/bos_inst/var		jfs2	May 17 13:57	rw,log=/dev/hd8
/dev/bos_hd10opt	/bos_inst/opt		jfs2	May 17 13:57	rw,log=/dev/hd8
/usr/lib	/bos_inst/usr/lib/multibos_chroot/usr/lib		namefs	May 17 13:57	rw
/usr/ccs/lib	/bos_inst/usr/lib/multibos_chroot/usr/ccs/lib		namefs	May 17 13:57	rw
/tmp	/bos_inst/tmp		namefs	May 17 13:57	rw

```
MULTIBOS> exit
Script done, file is /etc/multibos/logs/scriptlog.080517135722.txt
Stopping multibos root shell ...
Compressing script log file ...
Compressed script log file is /etc/multibos/logs/scriptlog.080517135722.txt.Z

+-----+
Mount Processing
+-----+

Unmounting all standby BOS file systems ...
Unmounting /bos_inst/opt
Unmounting /bos_inst/var
Unmounting /bos_inst/usr
Unmounting /bos_inst

Log file is /etc/multibos/logs/op.alog
Return Status = SUCCESS
```

```
bash-3.00# df -g
```

Filesystem	GB	blocks	Free	%Used	Iused	%Iused	Mounted on
/dev/hd4	1.00		0.96	4%	2779	2%	/
/dev/hd2	4.00		2.12	47%	40760	8%	/usr
/dev/hd9var	4.00		3.98	1%	449	1%	/var
/dev/hd3	4.00		3.99	1%	95	1%	/tmp
/dev/hd1	1.00		1.00	1%	5	1%	/home
/proc	-		-	-	-	-	/proc
/dev/hd10opt	1.00		0.92	8%	2249	2%	/opt
/dev/swdump	15.00		3.15	80%	5308	1%	/swdump
/dev/nim1	6.00		3.77	38%	5	1%	/nim1

Start up selection between the multibos Environments.

In the SMS utility, carefully look at the boot options.

```
PowerPC Firmware
Version SF240_338
SMS 1.6 (c) Copyright IBM Corp. 2000,2005 All rights reserved.
-----
Select Device
Device   Current   Device
Number  Position  Name
1.       -         Ethernet
          ( loc=U5791.001.99B0TG6-P1-C01-T1 )
2.       -         Ethernet
          ( loc=U5791.001.99B0TG6-P1-C01-T2 )
3.       2         SCSI 146814 MB Harddisk, part=2 (AIX 5.3.0)
          ( loc=U5791.001.99B0TG6-P1-T5-L8-L0 )
4.       1         SCSI 146814 MB Harddisk, part=4 (AIX 5.3.0)
          ( loc=U5791.001.99B0TG6-P1-T5-L8-L0 )
```

Same disk bootable entries twice displayed. One of them have the option as part=2 and another one as part=4. They are the blv partitions of active BOS and standby BOS.

```
Saving Base Customize Data to boot disk
Starting the sync daemon
Starting the error daemon
System initialization completed.
Setting tunable parameters...complete
Starting Multi-user Initialization
  Performing auto-varyon of Volume Groups
  Activating all paging spaces
swapon: Paging device /dev/hd6 is already active.

The current volume is: /dev/hd1
Primary superblock is valid.

The current volume is: /dev/bos_hd10opt
Primary superblock is valid.
  Performing all automatic mounts
Multi-user initialization completed
Initializing multibos methods ...
Initializing log /etc/multibos/logs/op.alog ...
Gathering system information ...

Log file is /etc/multibos/logs/op.alog
```

```

# hostname
localhost
# df -g
Filesystem      GB blocks      Free %Used      Iused %Iused Mounted on
/dev/bos_hd4    1.00           0.96   4%       2779    2% /
/dev/bos_hd2    4.00           2.12  47%      40760   8% /usr
/dev/bos_hd9var 4.00           3.98   1%        449    1% /var
/dev/hd3        4.00           3.99   1%         96    1% /tmp
/dev/hd1        1.00           1.00   1%          5    1% /home
/proc           -              -      -          -     - /proc
/dev/bos_hd10opt 1.00           0.92   8%       2249   2% /opt
# ifconfig -a
en0: flags=5e080863,c0<UP,BROADCAST,NOTRAILERS,RUNNING,SIMPLEX,MULTICAST,GROUPRT,6
D,CHAIN>
    inet 10.10.10.21 netmask 0xffffffff broadcast 10.10.10.255
    tcp_sendspace 131072 tcp_recvspace 65536
lo0: flags=e08084b<UP,BROADCAST,LOOPBACK,RUNNING,SIMPLEX,MULTICAST,GROUPRT,64BIT>
    inet 127.0.0.1 netmask 0xff000000 broadcast 127.255.255.255
    inet6 ::1/0
    tcp_sendspace 131072 tcp_recvspace 131072 rfc1323 1
# lsvg -l rootvg
rootvg:
LV NAME          TYPE          LPs   PPs   PVs  LV STATE      MOUNT POINT
hd5              boot          1     1     1    closed/syncd  N/A
hd6              paging        16    16    1    open/syncd    N/A
hd8              jfs2log       1     1     1    open/syncd    N/A
hd4              jfs2          4     4     1    closed/syncd  /bos_inst
hd2              jfs2          16    16    1    closed/syncd  /bos_inst/usr
hd9var           jfs2          16    16    1    closed/syncd  /bos_inst/var
hd3              jfs2          16    16    1    open/syncd    /tmp
hd1              jfs2          4     4     1    open/syncd    /home
hd10opt          jfs2          4     4     1    closed/syncd  /bos_inst/opt
bos_hd5          boot          1     1     1    closed/syncd  N/A
bos_hd4          jfs2          4     4     1    open/syncd    /
bos_hd2          jfs2          16    16    1    open/syncd    /usr
bos_hd9var       jfs2          16    16    1    open/syncd    /var
bos_hd10opt      jfs2          4     4     1    open/syncd    /opt
# █

```

```

bash-3.00# multibos -mX
Initializing multibos methods ...
Initializing log /etc/multibos/logs/op.alog ...
Gathering system information ...

+-----+
BOS Mount Operation
+-----+
Verifying operation parameters ...

+-----+
Mount Processing
+-----+
Mounting all standby BOS file systems ...
Mounting /bos_inst
Mounting /bos_inst/usr
Mounting /bos_inst/var
Mounting /bos_inst/opt

Log file is /etc/multibos/logs/op.alog
Return Status = SUCCESS
bash-3.00# df -g

```

Filesystem	GB	blocks	Free	%Used	Iused	%Iused	Mounted on
/dev/bos_hd4	1.00		0.96	4%	2780	2%	/
/dev/bos_hd2	4.00		2.12	47%	40760	8%	/usr
/dev/bos_hd9var	4.00		3.98	1%	449	1%	/var
/dev/hd3	4.00		3.99	1%	96	1%	/tmp
/dev/hd1	1.00		1.00	1%	5	1%	/home
/proc	-		-	-	-	-	/proc
/dev/bos_hd10opt	1.00		0.92	8%	2249	2%	/opt
/dev/hd4	1.00		0.96	4%	2758	2%	/bos_inst
/dev/hd2	4.00		2.12	47%	40760	8%	/bos_inst/usr
/dev/hd9var	4.00		3.98	1%	443	1%	/bos_inst/var
/dev/hd10opt	1.00		0.92	8%	2249	2%	/bos_inst/opt

Removing the standby BOS

```

bash-3.00# multibos -RX
Initializing multibos methods ...
Initializing log /etc/multibos/logs/op.alog ...
Gathering system information ...

+-----+
Remove Operation
+-----+
Verifying operation parameters ...

+-----+
Boot Partition Processing
+-----+
Active boot logical volume is bos_hd5.
Standby boot logical volume is hd5.

+-----+
Mount Processing
+-----+
Unmounting all standby BOS file systems ...
Unmounting /bos_inst/opt
Unmounting /bos_inst/var
Unmounting /bos_inst/usr
Unmounting /bos_inst

+-----+
File Systems
+-----+
Removing all standby BOS file systems ...
Removing standby BOS file system /bos_inst/opt
Removing standby BOS file system /bos_inst/var
Removing standby BOS file system /bos_inst/usr
Removing standby BOS file system /bos_inst

```

```

+-----+
Logical Volumes
+-----+
Removing all standby BOS logical volumes ...
Removing standby BOS logical volume hd5

+-----+
Bootlist Processing
+-----+
Verifying operation parameters ...
Setting bootlist to logical volume bos_hd5 on hdisk0.
ATTENTION: firmware recovery string for BLV ():
boot

Log file is /etc/multibos/logs/op.alog
Return Status = SUCCESS

```

```

bash-3.00# lsvg -l rootvg
rootvg:
LV NAME          TYPE      LPs    PPs    PVs    LV STATE    MOUNT POINT
hd6              paging    16     16     1      open/syncd  N/A
hd8              jfs2log   1       1      1      open/syncd  N/A
hd3              jfs2     16     16     1      open/syncd  /tmp
hd1              jfs2      4       4      1      open/syncd  /home
bos_hd5          boot      1       1      1      closed/syncd N/A
bos_hd4          jfs2      4       4      1      open/syncd  /
bos_hd2          jfs2     16     16     1      open/syncd  /usr
bos_hd9var       jfs2     16     16     1      open/syncd  /var
bos_hd10opt      jfs2      4       4      1      open/syncd  /opt

```

```

bash-3.00# df -g
Filesystem      GB blocks      Free %Used      Iused %Iused Mounted on
/dev/bos_hd4    1.00          0.96   4%      2765    2% /
/dev/bos_hd2    4.00          2.12  47%     40760   8% /usr
/dev/bos_hd9var 4.00          3.98   1%       445    1% /var
/dev/hd3        4.00          3.99   1%        96    1% /tmp
/dev/hd1        1.00          1.00   1%         5    1% /home
/proc           -             -       -         -     - /proc
/dev/bos_hd10opt 1.00          0.92   8%      2249    2% /opt

```

PowerPC Firmware

Version SF240_338

SMS 1.6 (c) Copyright IBM Corp. 2000,2005 All rights reserved.

Select Device

Device Number	Current Position	Device Name
1.	-	Ethernet (loc=U5791.001.99B0TG6-P1-C01-T1)
2.	-	Ethernet (loc=U5791.001.99B0TG6-P1-C01-T2)
3.	1	SCSI 146814 MB Harddisk, part=4 (AIX 5.3.0) (loc=U5791.001.99B0TG6-P1-T5-L8-L0)

Some more

```
# hostname
localhost
# df -g
Filesystem      GB blocks      Free %Used      Iused %Iused Mounted on
/dev/bos_hd4    1.00          0.96   4%       2763    2% /
/dev/bos_hd2    4.00          2.12  47%      40760   8% /usr
/dev/bos_hd9var 4.00          3.98   1%        444    1% /var
/dev/hd3        4.00          3.99   1%         96    1% /tmp
/dev/hd1        1.00          1.00   1%         5     1% /home
/proc           -             -       -          -       - /proc
/dev/bos_hd10opt 1.00          0.92   8%       2249   2% /opt
# lsvg -l rootvg
rootvg:
LV NAME          TYPE          LPs   PPs   PVs   LV STATE      MOUNT POINT
hd6              paging       16    16    1     open/syncd    N/A
hd8              jfs2log      1     1     1     open/syncd    N/A
hd3              jfs2         16    16    1     open/syncd    /tmp
hd1              jfs2         4     4     1     open/syncd    /home
bos_hd5          boot         1     1     1     closed/syncd  N/A
bos_hd4          jfs2         4     4     1     open/syncd    /
bos_hd2          jfs2        16    16    1     open/syncd    /usr
bos_hd9var       jfs2        16    16    1     open/syncd    /var
bos_hd10opt      jfs2         4     4     1     open/syncd    /opt
```

```

# multibos -sX
Initializing multibos methods ...
Initializing log /etc/multibos/logs/op.aalog ...
Gathering system information ...

+-----+
Setup Operation
+-----+
Verifying operation parameters ...
Creating image.data file ...

+-----+
Logical Volumes
+-----+
Creating standby BOS logical volume hd5
Creating standby BOS logical volume hd4
Creating standby BOS logical volume hd2
Creating standby BOS logical volume hd9var
Creating standby BOS logical volume hd10opt

+-----+
File Systems
+-----+
Creating all standby BOS file systems ...
Creating standby BOS file system /bos_inst on logical volume hd4
Creating standby BOS file system /bos_inst/usr on logical volume hd2
Creating standby BOS file system /bos_inst/var on logical volume hd9var
Creating standby BOS file system /bos_inst/opt on logical volume hd10opt

+-----+
Mount Processing
+-----+
Mounting all standby BOS file systems ...
Mounting /bos_inst
Mounting /bos_inst/usr
Mounting /bos_inst/var
Mounting /bos_inst/opt

```

```

+-----+
BOS Files
+-----+
Including files for file system /
Including files for file system /usr
Including files for file system /var
Including files for file system /opt

Copying files using backup/restore utilities ...
Percentage of files copied: 0.00%
Percentage of files copied: 1.66%
Percentage of files copied: 3.33%
Percentage of files copied: 5.00%
Percentage of files copied: 6.67%
Percentage of files copied: 8.34%

```

```

# lsvg -l rootvg
rootvg:
LV NAME          TYPE      LPs    PPs    PVs    LV STATE      MOUNT POINT
hd5              boot      1      1      1      closed/syncd  N/A
hd6              paging    16     16     1      open/syncd    N/A
hd8              jfs2log   1      1      1      open/syncd    N/A
hd4              jfs2      4      4      1      open/syncd    /bos_inst
hd2              jfs2      16     16     1      open/syncd    /bos_inst/usr
hd9var           jfs2      16     16     1      open/syncd    /bos_inst/var
hd3              jfs2      16     16     1      open/syncd    /tmp
hd1              jfs2      4      4      1      open/syncd    /home
hd10opt          jfs2      4      4      1      open/syncd    /bos_inst/opt
bos_hd5          boot      1      1      1      closed/syncd  N/A
bos_hd4          jfs2      4      4      1      open/syncd    /
bos_hd2          jfs2      16     16     1      open/syncd    /usr
bos_hd9var       jfs2      16     16     1      open/syncd    /var
bos_hd10opt      jfs2      4      4      1      open/syncd    /opt
# df -g
Filesystem      GB blocks      Free %Used      Iused %Iused Mounted on
/dev/bos_hd4    1.00           0.96  4%          2784   2% /
/dev/bos_hd2    4.00           2.12  47%         40760  8% /usr
/dev/bos_hd9var 4.00           3.98  1%           449   1% /var
/dev/hd3        4.00           3.99  1%           111   1% /tmp
/dev/hd1        1.00           1.00  1%            5   1% /home
/proc           -              -     -            -     - /proc
/dev/bos_hd10opt 1.00           0.92  8%          2249   2% /opt
/dev/hd4        1.00           0.96  4%          2758   2% /bos_inst
/dev/hd2        4.00           3.99  1%           224   1% /bos_inst/usr
/dev/hd9var     4.00           4.00  1%            4   1% /bos_inst/var
/dev/hd10opt    1.00           0.92  8%          2249   2% /bos_inst/opt

```

```

root 483340 446694 2 14:45:40 pts/0 0:00 backbyname -i -q -f -
root 487654 532562 0 14:47:57 pts/1 0:00 more
root 491622 409814 0 14:45:40 pts/0 0:00 /usr/bin/ksh /usr/sbin/multibos -sX
root 495862 86180 0 14:44:29 - 0:00 /usr/sbin/rsct/bin/IBM.HostRMd
root 499956 86180 0 14:44:27 - 0:00 /usr/sbin/rsct/bin/IBM.ERrmd
root 504056 86180 0 14:44:27 - 0:00 /usr/sbin/rsct/bin/IBM.AuditRMd
root 508020 360680 0 14:47:17 - 0:00 telnetd -a
root 512252 516350 0 14:44:34 pts/0 0:00 -ksh
root 516350 360680 0 14:44:34 - 0:00 telnetd -a
root 520232 491622 2 14:45:40 pts/0 0:01 restbyname -q -f -

```

```

Percentage of files copied: 95.08%
Percentage of files copied: 96.75%
Percentage of files copied: 98.42%
Percentage of files copied: 100.00%

```

```

+-----+
Boot Partition Processing
+-----+
Active boot logical volume is bos_hd5.
Standby boot logical volume is hd5.
Creating standby BOS boot image on boot logical volume hd5
bosboot: Boot image is 33193 512 byte blocks.

+-----+
Mount Processing
+-----+
Unmounting all standby BOS file systems ...
Unmounting /bos_inst/opt
Unmounting /bos_inst/var
Unmounting /bos_inst/usr
Unmounting /bos_inst

+-----+
Bootlist Processing
+-----+
Verifying operation parameters ...
Setting bootlist to logical volume hd5 on hdisk0.
ATTENTION: firmware recovery string for BLV ():
boot

Log file is /etc/multibos/logs/op.alog
Return Status = SUCCESS

```

```

# multibos -mX
Initializing multibos methods ...
Initializing log /etc/multibos/logs/op.alog ...
Gathering system information ...

+-----+
BOS Mount Operation
+-----+
Verifying operation parameters ...

+-----+
Mount Processing
+-----+
Mounting all standby BOS file systems ...
Mounting /bos_inst
Mounting /bos_inst/usr
Mounting /bos_inst/var
Mounting /bos_inst/opt

Log file is /etc/multibos/logs/op.alog
Return Status = SUCCESS

```

```
# df -g
Filesystem      GB blocks      Free %Used      Iused %Iused Mounted on
/dev/bos_hd4    1.00           0.96   4%       2784    2% /
/dev/bos_hd2    4.00           2.12  47%      40760   8% /usr
/dev/bos_hd9var 4.00           3.98   1%        449    1% /var
/dev/hd3        4.00           3.99   1%         96    1% /tmp
/dev/hd1        1.00           1.00   1%          5    1% /home
/proc           -              -      -          -     - /proc
/dev/bos_hd10opt 1.00           0.92   8%       2249   2% /opt
/dev/hd4        1.00           0.96   4%       2759   2% /bos_inst
/dev/hd2        4.00           2.12  47%      40760   8% /bos_inst/usr
/dev/hd9var     4.00           3.98   1%         438    1% /bos_inst/var
/dev/hd10opt    1.00           0.92   8%       2249   2% /bos_inst/opt
```

```
bash-3.00# bootlist -m normal -o
hdisk0 blv=hd5
hdisk0 blv=bos_hd5
```

After reboot....

```
# hostname
localhost
# df -g
Filesystem      GB blocks      Free %Used      Iused %Iused Mounted on
/dev/hd4        1.00           0.96   4%       2777   2% /
/dev/hd2        4.00           2.12  47%      40760   8% /usr
/dev/hd9var     4.00           3.98   1%         449    1% /var
/dev/hd3        4.00           3.99   1%         95    1% /tmp
/dev/hd1        1.00           1.00   1%          5    1% /home
/proc           -              -      -          -     - /proc
/dev/hd10opt    1.00           0.92   8%       2249   2% /opt
```

```
# lsvg -l rootvg
rootvg:
LV NAME          TYPE          LPs   PPs   PVs   LV STATE      MOUNT POINT
hd5              boot          1     1     1     closed/syncd  N/A
hd6              paging        16    16    1     open/syncd    N/A
hd8              jfs2log       1     1     1     open/syncd    N/A
hd4              jfs2          4     4     1     open/syncd    /
hd2              jfs2          16    16    1     open/syncd    /usr
hd9var           jfs2          16    16    1     open/syncd    /var
hd3              jfs2          16    16    1     open/syncd    /tmp
hd1              jfs2          4     4     1     open/syncd    /home
hd10opt          jfs2          4     4     1     open/syncd    /opt
bos_hd5          boot          1     1     1     closed/syncd  N/A
bos_hd4          jfs2          4     4     1     closed/syncd  /bos_inst
bos_hd2          jfs2          16    16    1     closed/syncd  /bos_inst/usr
bos_hd9var       jfs2          16    16    1     closed/syncd  /bos_inst/var
bos_hd10opt      jfs2          4     4     1     closed/syncd  /bos_inst/opt
```

```

# multibos -RX
Initializing multibos methods ...
Initializing log /etc/multibos/logs/op.alog ...
Gathering system information ...

+-----+
Remove Operation
+-----+
Verifying operation parameters ...

+-----+
Boot Partition Processing
+-----+
Active boot logical volume is hd5.
Standby boot logical volume is bos_hd5.

+-----+
Mount Processing
+-----+
Unmounting all standby BOS file systems ...

+-----+
File Systems
+-----+
Removing all standby BOS file systems ...
Removing standby BOS file system /bos_inst/opt
Removing standby BOS file system /bos_inst/var
Removing standby BOS file system /bos_inst/usr
Removing standby BOS file system /bos_inst

```

```

+-----+
Logical Volumes
+-----+
Removing all standby BOS logical volumes ...
Removing standby BOS logical volume bos_hd5

+-----+
Bootlist Processing
+-----+
Verifying operation parameters ...
Setting bootlist to logical volume hd5 on hdisk0.
ATTENTION: firmware recovery string for BLV ():
boot

Log file is /etc/multibos/logs/op.alog
Return Status = SUCCESS

```

We are back from where we started with a plain AIX instance.

```
bash-3.00# which multibos
/usr/sbin/multibos
bash-3.00# lslpp -w /usr/sbin/multibos
```

File	Fileset	Type
/usr/sbin/multibos	bos.rte.bosinst	File

```
# bash-3.00# lslpp -f bos.rte.bosinst | grep multi
roc /usr/sbin/multibos
LV /usr/lpp/bos/README.multibos MOUNT POINT
hd5 boot 1 1 1 closed/syncd N/A
hd6 paging 16 16 1 open/syncd N/A
hd8 jfs2log 1 1 1 open/syncd N/A
hd4 jfs2 4 4 1 open/syncd /
hd2 jfs2 16 16 1 open/syncd /usr
hd9var jfs2 16 16 1 open/syncd /var
hd3 jfs2 16 16 1 open/syncd /tmp
hd1 jfs2 4 4 1 open/syncd /home
hd10opt jfs2 4 4 1 open/syncd /opt
# df -g
```

Filesystem	GB blocks	Free	%Used	Iused	%Iused	Mounted on
/dev/hd4	1.00	0.96	4%	2763	2%	/
/dev/hd2	4.00	2.12	47%	40760	8%	/usr
/dev/hd9var	4.00	3.98	1%	445	1%	/var
/dev/hd3	4.00	3.99	1%	95	1%	/tmp
/dev/hd1	1.00	1.00	1%	5	1%	/home
/proc	-	-	-	-	-	/proc
/dev/hd10opt	1.00	0.92	8%	2249	2%	/opt

```
bash-3.00# bootlist -m normal -o -v
'ibm,max-boot-devices' = 0x5
NVRAM variable: (boot-device=/pci@800000020000012/pci@2,6/scsi@1/sd@8:2)
Path name: (/pci@800000020000012/pci@2,6/scsi@1/sd@8:2)
match_specific_info: ut=disk/scsi/scsd
hdisk0 blv=hd5
```

UNDOCUMENTED OPERATIONS

In addition to operations documented in the man page and publications, an undocumented verify operation is run from the inittab during boot. The inittab entry looks as such:

```
mbverify:23456789:wait:/usr/sbin/multibos -V 2>&1 | alog -t boot > /dev/console
```

It is highly recommended that the user not modify this entry. This verify operation allows the multibos utility to synchronize changes in logical volumes and filesystems between the active and standby instances. This entry also synchronizes the ODM and devices on initial boot after a mksysb restore. Without this operation, both the active and standby instances could become inconsistent with normal filesystem and logical volume operations.

The verify operation only synchronizes data when the multibos utility recognizes that the system has changed the instance from which it is booting or that it is the first boot after a backup is restored. Therefore, it should be of little impact for consecutive boots from the same instance.

If this inittab entry is blocked, removed or modified, the user is responsible for synchronization of the instances.

This inittab entry (if runnable) is removed upon removal of the standby BOS (multibos -R).

SUPPORTED METHODS OF BACKUP/RECOVERY

Currently, the only supported method of backup and recovery of a rootvg with multiple instances is mkysyb through CD, tape, or NIM. If the standby multibos instance has been removed and the active instance uses the traditional LV names (eg, no bos_ tag appended to the LV), other methods of backup and recovery of the rootvg are expected to work.

If the standby BOS was mounted during the creation of the mkysyb backup, it will be restored and synchronized on the first boot from the restored data. However, if the standby BOS was not mounted during the creation of the mkysyb backup, the synchronization on reboot will remove the unusable standby BOS.

NOTE: No Alternate Disk Installation operations are currently supported for this environment. The user should not confuse alt_disk_install mkysyb or alt_mkysyb with the above supported mkysyb methods.

NOTES REGARDING DUAL BOOT

To keep track of the boot logical volume (BLV) that maps to the running boot image, the new device /dev/ipl_blv is a link to the active BLV device. The bosboot and rc.boot files manage the /dev/ipl_blv link. The /dev/ipldevice link still maps to the disk containing the boot image.

By default, all operations map to the active BOS. This should provide for compatibility with all non-multibos user level operations.

The chpv -c operation clears the boot record of the disk. Using this operation will clear mapping to both the active and standby BLVs. Care should be taken when using chpv -c.

A bosboot operation on the active BLV will initialize the boot record mapping for the active BOS. In addition, the multibos bosboot operation (multibos -B) may be used to initialize the boot record mapping to the standby BOS.

The -t option is used to block any change of the bootlist by the multibos utility.

MULTIBOS DATA FILES

All log files are kept in the /etc/multibos/logs directory. The following are examples of files that may be found in this directory:

- op.alog : A circular alog file of all multibos operations.
- scriptlog.<timestamp>.txt : A log of commands being run during the current shell operation.
- scriptlog.<timestamp>.txt.Z : A compressed log of commands run during a previous shell operation.

In addition, the bootlog contains redundant logging of all multibos operations that occur during boot (eg, the verify that attempts synchronization from inittab).

Multibos private data is stored in the /etc/multibos/data directory, the logical volume control block (LVCB) of logical volumes that were the source or target of a copy, and the /etc/filesystems entries that were the source or target of a copy. The following are examples of files found in the /etc/multibos/data directory:

- acttag : The multibos tag for the active BOS.
- sbyfslist : The list of filesystems private to the standby BOS.
- sbylvlist : The list of logical volumes private to the standby BOS.
- sbytag : The multibos tag for the standby BOS.

The following may be seen in the fs field of the logical volumes that were the source or target of a copy:

```
mb=<TAG>:mbverify=<integer>
```

The following may be seen in /etc/filesystems as an attribute of filesystems that were the source or target of a copy:

```
mb = <TAG>
```

The user should not modify multibos private data.

To prevent multibos operations from working simultaneously, the directory /etc/multibos/locks contains lock files. The following is an example of a file that may be found in this directory:

- global_lock : The process ID (PID) of the currently running multibos operation.

If a multibos operation exited unexpectedly and was not able to clean up, it may be necessary to remove this file. The user should check that the PID is not running before removing this file.

Dlpar concordance

DLPAR is the facility on the IBM POWER4, POWER5 & POWER6 hardware that allows users to dynamically move PCI slots, memory and processors among LPARs, increasing and decreasing resources on the fly. DLPAR is commonly used to move CD drives, DVD-RAM drives and tape drives between LPARs, rather than attaching one to each LPAR. Since its initial introduction with POWER4 and AIX V5.2, DLPAR has enhanced its capabilities and plays well within the new shared environment in System P machines. DLPAR requires HMC at certain minimum levels and also requires that the OS be at certain levels.

Dlpar generally gets activated automatically on an POWER5/POWER6 lpar just 5-10minutes after the HMC is able to communicate with the AIX rmc subsystem through the HMC's open network, where "lpar communication" is enabled.

Below are the brief steps to check the rmc subsystem connectivity between HMC and AIX running on the lpar.

1. Ping the ip-address of the AIX in the lpar from the HMC.
2. Ping the hmc public (open) network ip-address from AIX ip-address.
Enable AIX lpar communication option on this interface.
3. Ensure that in the AIX no system filesystems are 100%.
4. Check in the HMC using the following commands.

```
hscroot@inghmc3:~> lspartition -dlpar
<#0> Partition:<3*9119-590*8328F20, , 10.10.10.24>
      Active:<1>, OS:<AIX, 5.3>, DCaps:<0x2f>, CmdCaps:<0xb, 0xb>, PinnedMem:<285>
<#1> Partition:<2*9119-590*8328F20, , 10.10.10.23>
      Active:<0>, OS:<AIX, 5.3>, DCaps:<0x0>, CmdCaps:<0x0, 0xb>, PinnedMem:<337>
<#2> Partition:<1*9119-590*8328F20, , 10.10.10.22>
      Active:<1>, OS:<AIX, 5.3>, DCaps:<0x2f>, CmdCaps:<0xb, 0xb>, PinnedMem:<924>
<#3> Partition:<5*9119-590*8328F20, , 10.10.10.21>
      Active:<0>, OS:<AIX, 5.3>, DCaps:<0x0>, CmdCaps:<0x0, 0xb>, PinnedMem:<480>
<#4> Partition:<4*9119-590*8328F20, , 10.10.10.25>
      Active:<1>, OS:<AIX, 5.3>, DCaps:<0x2f>, CmdCaps:<0xb, 0xb>, PinnedMem:<386>
```

```
hscroot@localhost:~> lspartition -dlpar
<#0> Partition:<3*9119-590*8328F20, , 10.10.10.24>
      Active:<1>, OS:<AIX, 5.3>, DCaps:<0x2f>, CmdCaps:<0xb, 0xb>, PinnedMem:<392>
<#1> Partition:<2*9119-590*8328F20, , 10.10.10.23>
      Active:<1>, OS:<AIX, 5.3>, DCaps:<0x2f>, CmdCaps:<0xb, 0xb>, PinnedMem:<268>
<#2> Partition:<1*9119-590*8328F20, , 10.10.10.22>
      Active:<1>, OS:<AIX, 5.3>, DCaps:<0x2f>, CmdCaps:<0xb, 0xb>, PinnedMem:<778>
<#3> Partition:<5*9119-590*8328F20, , 10.10.10.21>
      Active:<1>, OS:<AIX, 5.3>, DCaps:<0x2f>, CmdCaps:<0xb, 0xb>, PinnedMem:<469>
<#4> Partition:<4*9119-590*8328F20, , 10.10.10.25>
      Active:<1>, OS:<AIX, 5.3>, DCaps:<0x2f>, CmdCaps:<0xb, 0xb>, PinnedMem:<313>
```

Look for the Variable "Active" and Dcaps:

Active:<1> dlpar OK

Active:<0> dlpar failure

DCaps:<0x2f> dlpdr OK
 DCaps:<0x0> dlpdr failure

```
hscroot@inghmc3:~> lspartition -sfp
<#0> Partition:<4, , 10.10.10.25>
      Active:<3>, OS:<AIX, 5.3>
<#1> Partition:<3, , 10.10.10.24>
      Active:<3>, OS:<AIX, 5.3>
<#2> Partition:<2, , 10.10.10.23>
      Active:<0>, OS:<AIX, 5.3>
<#3> Partition:<1, , 10.10.10.22>
      Active:<3>, OS:<AIX, 5.3>
<#4> Partition:<5, , 10.10.10.21>
      Active:<0>, OS:<AIX, 5.3>
```

5. Check in the AIX using the following commands.

```
# lssrc -a | grep rsct
ctrmc          rsct          438284      active
IBM.CSMAgentRM rsct_rm        626726      active
IBM.ERRM       rsct_rm        389274      active
IBM.HostRM     rsct_rm        479240      active
IBM.ServiceRM  rsct_rm        463096      active
IBM.AuditRM    rsct_rm        454658      active
IBM.DRM        rsct_rm        327856      active
IBM.LPRM       rsct_rm        499748      active
ctcas          rsct           inoperative
```

ctrmc: is a RMC subsystem

ctcas: is for security verification

IBM.CSMAgentRM: is for handshaking between the partition and HMC

IBM.HostRM: is for obtaining OS information.

I IBM.DRM: is for executing the DLPAR command on the partition.

IBM.DMSRM: is for tracking statuses of partitions

IBM.LparCmdRM: is for DLPAR operation on HMC

```
# pwd
/usr/sbin/rsct/bin
# ./lssrc "IBM.ManagementServer"
Resource Persistent Attributes for IBM.ManagementServer
resource 1:
  Name           = "10.10.10.12"   <==== HMC ip-address.
  Hostname       = "10.10.10.12"
  ManagerType    = "HMC"
  LocalHostname  = "10.10.10.21"
  ClusterTM      = "9078-160"
  ClusterSNum    = ""
  ActivePeerDomain = ""
  NodeNameList   = {"standbysvr"}
```

6. If rmct subsystems are not able to establish connectivity then do the following.

- a) Ensure that AIX and HMC are able to ping each other.
- b) Reboot the HMC, this ensures that the RMC subsystems get restarted in the HMC.
- c) Run the following commands from AIX and wait for 10 -15 minutes.

```
./rmcctrl -K  
./rmcctrl -A
```

The DiagnoseHMC is installed in the /opt/csm/csm/bin/diagnostics directory on both AIX and HMC. On AIX, it is included in the csm.core fileset, which is installed by default. On HMC, the command requires the root authority.

POWER4 systems also allow dlpdr, but some manual configuration is must to get it established. The prime criteria is that name resolution must be there for the dlpdr to work. In simple terms, hmc should have a host entry for the AIX lpars and AIX lpars should have host entry for the HMC. DNS also can also be employed. But prefer to go with the host entries. Then `./rmcctrl -K` followed by `./rmcctrl -A`, wait for 5-10minutes and then check for the dlpdr option in the HMC.