AIX 5L Version 5.3

Release Notes

SC23-5201-02

AIX 5L Version 5.3

Release Notes

SC23-5201-02

Note

Before using this information and the product it supports, read the information in Appendix E, “Notices,” on page 69.

Fifth Edition (July 2006)

Contents

Chapter 1. Read this before installation

Note: This software may contain errors that could result in critical business impact. It is highly recommended that you install the latest available fixes prior to using this software. Fixes can be obtained from IBM System p support at the following Web site:

`http://www.ibm.com/servers/eserver/support/unixservers/aixfixes.html`

These Release Notes support AIX 5L™ Version 5.3 with the recommended 5300-05 Technology Level package.

The AIX 5L Version 5.3 Release Notes include information that helps you install AIX 5L Version 5.3. To view the most current version of the Release Notes, go to the online Release Notes in the **System p and AIX Information Center** where the latest changes in the English version are marked with a ">" in the left margin. The information center is located at the following Web site:

`Release Notes Index for AIX 5.3 and Expansion Pack`

(http://publib.boulder.ibm.com/infocenter/pseries/v5r3/index.jsp?topic=/com.ibm.aix.resources/53relnotes.htm)

Installation tips

The latest installation hints and tips are available from the IBM® Subscription Service for UNIX® servers at:

`http://www14.software.ibm.com/webapp/set2/subscriptions/pqvcmjd`

These tips might contain information that are critical for successful information of this software.

Software License Agreements (SLA)

There are instances where the Software Licence Agreements might not be displayed correctly. In this event, the License Agreements can be viewed in all languages at the following Web site:

`http://www.ibm.com/software/sla/sladb.nsf`

What's new in AIX

Read about the latest updates to the AIX 5L Version 5.3 operating system. *What's New in AIX* is located at the following Web site:

`http://publib.boulder.ibm.com/infocenter/pseries/v5r3/index.jsp`

Select **Other AIX resources** → **Other resources** → **What's new in AIX** (under **Related links**)

Service

Fixes and problem-solving databases

You can download AIX® fixes and search technical databases (including ″APARS″ and ″Tips for AIX administrators″), from the following IBM eServer™ Support Web site:

`http://www.ibm.com/servers/eserver/support/pseries/aixfixes.html`

Chapter 2. System requirements

Review the following information to determine the minimum and recommended system requirements needed to run AIX 5L Version 5.3.

Required hardware

Only Common Hardware Reference Platform (CHRP) machines are supported.

To see if you have a CHRP machine, log into the machine as the root user, and run the following command:

```
lscfg | grep Architecture
```

For more information about supported and unsupported items, see Appendix B, "AIX 5L Version 5.3 unsupported devices," on page 55.

Firmware

Some AIX systems might not boot from CD-ROM

Some AIX systems might not boot from CD-ROM because of firmware issues. To determine if your system will be affected by this problem, perform the following steps before you migrate or install a running AIX 4.3, AIX 5.1, or AIX 5.2 system:

1. At the command prompt, type the following:

   ```
   lscfg -vl cd*
   ```

2. Examine the data that is returned.

   If *Part Number* is 04N2964 and *ROS Level and ID* is less than or equal to `1_04` (for example, `1_02`, `1_01`, or `1_00`), contact your local service representative. Tell your service representative that your system requires the CD-ROM firmware upgrade that is described in RETAIN® TIP H1332.

   If the data returned does not match the data described in the preceding paragraph, your system is not affected by this problem.

Firmware upgrade required for System p 7025/7026 systems

The following System p models require an upgrade to their firmware in order to install and run AIX 5L Version 5.3:

- 7025-F80
- 7025-6F0
- 7025-6F1
- 7026-H80
- 7026-6H0
- 7026-6H1
- 7026-M80
- 7026-6M1

Note: You must upgrade the system firmware to the following levels before installing AIX 5L Version 5.3:

System	System level with support of AIX 5.3
7025-F80/6F0/6F1	CL040712 or later
7026-H80/6H0/6H1	CM040712 or later

System	System level with support of AIX 5.3
7026-M80/6M1	MM040712 or later

The required version of the firmware can be obtained from the following Web site:
`http://techsupport.services.ibm.com/server/mdownload`

Refer to the history section of the firmware level for the statement of AIX 5L Version 5.3 support.

Firmware upgrade required to support an alternate boot device

The requirement of upgrading firmware when installing or booting AIX 5.3 in the scenarios described in this section applies only to the systems listed in the table included at the end of this section.

Installation

Some systems might encounter installation problems when you are installing AIX 5.3 on a system that is currently installed with a earlier version of AIX or when you are installing an earlier version of AIX on a system currently installed with AIX 5.3.

A firmware upgrade is required to support the following installation scenarios:

1. When an installation device is specified using the SMS menus or the Open Firmware command-line *and* the specified device is not the first available device in the AIX bootlist.

   Note: SMS menus differ, depending on the system. On some systems, the installation device is specified under the Multiboot menus. On other systems, the installation device is specified under the Select Boot Options menus.
2. When the F5 key is selected during boot, which is the fast path to install from the CD device, *and* the CD device is not the first available device in the AIX bootlist.

No firmware upgrade is required if no installation device is specified in the SMS menus or the Open Firmware prompt *or* if the specified installation device is the first available device in the AIX bootlist.

Booting from an alternate boot device

An alternate boot device can be specified using the Open Firmware prompt. A firmware upgrade is required to support an alternate boot device (specified using the Open Firmware prompt) that is not also the first available device in the AIX bootlist.

The table below specifies which firmware levels are required to support these scenarios:

System	Microcode Level with AIX 5.3 fix
7013 or 7015 or 7017 -S70/S7A	20040716
7017-S80/S85	20040716
7025-F50	L04197
7025 or 7026-H50	L04197
7025 or 7026-H70	SST04195
7026-B80	NAN04194
7028-6C4/6E4	3R040323 or later
7028-6C1/6E1 and 9112-265	CLT04194 or later
7029-6C3/6E3 & 9114-275	3F041021 or later
7038-6M2	3K040323 or later
7039-651	3J040528 or later
7040-671/681	3H040528 or later*

System	Microcode Level with AIX 5.3 fix
7043-150/7046-B50	TCP04195
7043-260	SPX04197
7043 or 7044-270	SPH04194
7044-170	SPH04194
9076-260	SPX04197
9076-270	SPH04194
9076-N80	NI04195
9076-N81	NI04195
9076-WCN	L04197

* The 7040-671/681 Version 2 microcode does not support AIX 5.3.

The required version of the firmware can be obtained from the following Web site:
`http://techsupport.services.ibm.com/server/mdownload`

Storage adapter microcode

It is always important to update the adapter microcode to the latest fix level available. This is especially important for the following SCSI adapters:

- PCI-X Dual Channel Ultra320 SCSI Adapter (5712, 5710, 570B, 570A, 1974)
- PCI-X Dual Channel Ultra320 SCSI RAID Adapter (5703, 5711, 1975)
- Dual Channel SCSI RAID Enablement Card (5709, 5726, 1976)
- PCI-X Quad Channel U320 SCSI RAID Adapter (2780)
- PCI-XDDR Dual Channel Ultra320 SCSI Adapter (5736, 1912)
- PCI-XDDR Dual Channel U320 SCSI RAID Adapter (5737, 1913)
- Dual Channel SCSI RAID Enablement Card (5727, 5728, 1907)
- Dual Channel SCSI RAID Enablement Card (1908)

All these adapters support concurrent microcode download. Check for the latest adapter microcode updates at the following Web site:
`http://techsupport.services.ibm.com/server/mdownload`

Memory requirements

AIX 5L Version 5.3 minimum current memory requirements vary, based on the configuration.

Larger maximum memory configurations or additional devices will scale up the minimum current memory requirement. A general rule for a minimum current memory requirement for AIX 5L Version 5.3 is 256 MB–512 MB. A smaller minimum current memory requirement of 128 MB may support a configuration with a very small number of devices and where the maximum memory setting is set to match the current memory of 128 MB.

AIX 5L Version 5.3 requires the minimum current memory requirement to increase as the maximum memory configuration or the number of devices scales upward, or both.

Paging space requirements

AIX 5L Version 5.3 creates a 512 MB paging space (in the **/dev/hd6** directory) for all new and complete overwrite installations.

Disk requirements

AIX 5L Version 5.3 requires a minimum of 2.2 GB of physical disk space for the same set of installed filesets due to increased library sizes and additional function.

Note: The following measurements provide information about disk usage when you install AIX 5L Version 5.3 as compared to previous versions.

Base AIX Installation (Graphical System with CDE-Default)

Location	AIX 4.3.3 Allocated (Used)	AIX 5L for POWER™ Version 5.1 Allocated (Used)	AIX 5L Version 5.2 Allocated (Used)	AIX 5L Version 5.3 with the 5300-05 Recommended Technology Level Allocated (Used)
/	4 MB (2.5 MB)	8 MB (5.6 MB)	16 MB (10 MB)	16 MB (13 MB)
/usr	294 MB (279 MB)	385 MB (370 MB)	1040 MB (1034 MB)	1160 MB (1154 MB)
/var	4 MB (1.3 MB)	4 MB (1.4 MB)	16 MB (7 MB)	16 MB (7 MB)
/tmp	16 MB (0.6 MB)	20 MB (0.9 MB) (See note.)	32 MB (1.1 MB)	40 MB (1 MB)
/opt	N/A	4 MB (0.2 MB)	48 MB (26 MB)	64 MB (48 MB)

Note: If the **/tmp** directory has less than 32 MB, it is increased to 32 MB during a migration installation so that the AIX 5L Version 5.3 boot image is successfully created at the end of the migration.

During a migration installation, if **/opt** exists only as a directory and has less than 3 MB of data, then a new **/dev/hd10opt** logical volume and **/opt** file system are created, and the data that existed in the **/opt** directory is moved to the new **/opt** file system.

If there is more than 3 MB of data in the **/opt** directory, then the new logical volume and file system are not created.

If any existing file system has a mount point in the **/opt** directory, or a mount point of **/opt** itself, the new logical volume and file system are not created.

Supported devices

Parallel printer cable selection

The parallel printer cable must be changed to a cable that is IEEE1284-compliant if all of the following statements are true:

- Your system was manufactured after 1998.
- The printer is ″parallel attached.″
- The attached printer is not a dot-matrix printer.
- The output of the **lsdev -C -l ppa0** command contains the word `IEEE1284`.

If the output of the **lsdev** command contains the word Standard, or if the printer is a dot-matrix printer, an IEEE1284-compliant cable is not required.

Cables that are not IEEE1284-compliant may not correctly transmit data to high-speed printers. Loss of printer data may occur because the cables may not be capable of transmitting data at rates that are possible with newer ECP parallel ports.

Supported Enhanced Error Handling (EEH) devices

EEH is an I/O error detection, reporting, and recovery mechanism to increase system availability from such errors. In the current implementation, the EEH mechanism can recover I/O errors on the PCI bus for most devices. Information about the faulty component and nature of the error (recoverable versus permanent) is logged in the AIX error log.

For EEH to work, your system must have:

- AIX kernel support.
- AIX device driver support (dds). Most dds has full EEH recovery (with a few exceptions).
- EEH-capable hardware.
- Appropriate system firmware levels.

Certain hardware and firmware requirements must be met for EEH to work on a given system. Refer to your system guides to determine if EEH will work on your system.

Unsupported devices

For more information about unsupported devices, see "Unsupported EEH devices" in *Appendix A. AIX 5L Version 5.3 unsupported devices*.

Limitations and restrictions

Known limitations for POWER4 systems

Adapters

In Full System Partition mode, only one graphics adapter and USB adapter with one keyboard and mouse are allowed per system. Only one graphics adapter and USB adapter with one keyboard and mouse are allowed per logical partition, and a maximum of eight logical partitions that have a graphics adapter and USB adapter are allowed.

CPU Gard

Disable the CPU Gard functions if AIX 5L Version 5.3 and platform firmware levels older than October 2002 are used together by typing the following command:

```
chdev -l sys0 -a cpuguard='disable'
```

If platform firmware levels are upgraded, CPU Gard functions can be re-enabled by typing the following command:

```
chdev -l sys0 -a cpuguard='enable'
```

In either case, no system reboot is required for the changes to take effect.

System p™ 690 Memory

The System p 690 model 681 (7040-681) supports a maximum system memory size of 1 TB (terabyte) with appropriate memory Feature Codes installed.

AIX 5.3 and Linux® logical partitions can have nearly 512 GB logical partition memory sizes (about 503 GB after page table, POWER Hypervisor, and TCE table usage). AIX 5.3 and Linux logical partitions should have the **Small Real Mode Address Region** option selected on the HMC partition profile memory panel, and must be defined for logical partitions greater than 256 GB.

RAID capacity limitation

There are limits to the amount of disk drive capacity allowed in a single RAID array. Using the 32 bit kernel, there is a capacity limitation of 1 TB per RAID array. Using the 64 bit kernel, there is a capacity limitation of 2 TB per RAID array. For a RAID adapter and RAID enablement cards, this limitation is enforced by AIX when RAID arrays are created using the PCI-X SCSI Disk Array Manager. The adapters utilizing the PCI-X SCSI Disk Array Manager are:

- PCI-X Dual Channel Ultra320 SCSI RAID Adapter (5703, 5711, 1975)
- Dual Channel SCSI RAID Enablement Card (5709, 5726, 1976)
- PCI-X Quad Channel U320 SCSI RAID Adapter (2780)
- PCI-XDDR Dual Channel U320 SCSI RAID Adapter (5737, 1913)
- Dual Channel SCSI RAID Enablement Card (5727, 5728, 1907)
- Dual Channel SCSI RAID Enablement Card (1908)

When creating a RAID array up to 2 TB with Standalone Diagnostics, ensure that Version 5.3.0.40 or higher is used. Previous versions of the Standalone Diagnostics has a capacity limitation of 1 TB per RAID array.

InfiniBand limitation

There are two InfiniBand device drivers: one for the GX bus and the one for the PCIX bus. Both of these device drivers support only 64 bit kernel mode. Concurrent mode diagnostic support for the PCIX adapter is not provided.

Infiniband has a debug tracing capability built into the software stack for both adapters. The tracing is disabled by default. It can be controlled by the following commands:

- `IbDebugChk` - returns the current debug status
- `IbDebugOn` - enables the debug tracing
- `IbDebugOff` - disables the debug tracing

There is no tcpdump support for IPoIB. You can use the **ibstat** Command to display operational information.

ibstat command

The **ibstat** command displays Infiniband operational information pertaining to a specified Host Channel Adapter Device (HCAD). If an HCAD device name is not entered, status for all available HCAD's appear. You can display specific categories of information such as Node, Port, Interface, and Debug information, or you can choose to display all of the information categories.

You can use one of the following flags to narrow down your search results:

-d Displays current debug setting

-h Displays ibstat command usage

-i Displays Network Interface Information

-n Displays IB node information

-p Displays IB port information

-v Displays all IB device information

The following information appears on all valid calls and contains these fields:

Device Name
Displays the name of an available HCAD (example: iba0)

Port State
Displays the current state of each HCAD port

```
Down            Port is disabled.
Initialized     Port is enabled and issuing training sequences.
Armed           Port is trained and attempting to configure to the active state.
Active          Port is in a normal operational state.
Unknown         Port is in an invalid or unknown state.
```

If you specify an invalid Device_Name, the ibstat command produces error messages stating that it could not connect to the device such as:

```
IBSTAT: No device iba2 configured
IBSTAT: Device iba3 is not available.
```

Known problems

The following devices have limitations in the ability to update microcode with the microcode management feature:

- PCI 4 Channel Ultra3 SCSI RAID Adapter.
- CD-ROM and DVD-ROM Drives.
- RAID Devices.
- SSA devices and adapters.
- Inventory Scout does not properly handle some OEM adapters and devices.

For more information about these devices, see the readme files at the following Web site:

`http://techsupport.services.ibm.com/server/mdownload`

IBM 4.7 GB IDE Slimline DVD-RAM drive limitations

The following limitations apply to the IBM 4.7 GB IDE Slimline DVD-RAM drive:

- The DVD-RAM drive writes only to 4.7 GB and 9.4 GB DVD-RAM media and reads from CD-ROM, CD-R, CD-RW, DVD-ROM, and DVD-RAM media. If you try to write to CD media, you are prompted to insert DVD-R media.
- DVD video is not supported.
- Nonbootable mksysb backups fail. After you boot the system from the product media, the DVD-RAM does not mount to restore the mksysb backup.

Known problem writing to DVD drive

When creating system or volume group backups to Slimline or Virtual DVD-RAM drives you might see the following error:

```
/usr/bin/readcd: Invalid argument. Cannot send SCSI cmd via ioctl
 burn_cd: Command error.
```

You can verify the level of cdrecord by running the `lslpp -L cdrecord` command. It needs to be at least at level 1.9-6, which is available on this current AIX release or you can get it from:

`ftp://ftp.software.ibm.com/aix/freeSoftware/aixtoolbox/RPMS/ppc/cdrecord/cdrecord-1.9-6.aix5.2.ppc.rpm`

Limitation on placement of boot image on hard disk

The firmware in many of the PCI bus-based RS/6000® machines is limited in regard to the region of the hard disk from which it can read a boot image. This problem will not be encountered under most circumstances. The symptom of the problem is a failure to boot from hard disk, resulting in a message from firmware similar to `unrecognized Client Program format`.

Affected machines can be identified most easily as the machines that provide access to the firmware System Management Services by pressing the F1 key on the system-attached keyboard or the 1 key on a TTY keyboard.

Firmware on the affected machines cannot read the boot image from the hard disk if any part of the boot image is located beyond the 4 GB boundary on the hard disk. This is not a problem for most customers because the AIX installation process creates the boot logical volume at the beginning of the disk. This is achieved by using the **-a** flag with the **mklv** command and specifying **e** (which corresponds to **edge**) as the parameter for the **-a** flag. Using the **mklv** command with this parameter results in the boot logical volume being created at the edge of the hard disk, and the resulting address that the firmware uses to

read the boot image will be within a safe range. The AIX installation process has always created the boot logical volume near the edge of the hard disk because that region of the hard disk has the slowest access time, and this allows other regions of the hard disk to be used by file systems that can benefit from increased performance.

The only way that you can encounter this problem is by creating and initializing a new boot logical volume that extends past the 4 GB boundary of the hard disk.

In almost all cases, you do not need to create a new boot logical volume, but if you do, use the **lsvg** and **lslv** commands to verify that the newly created boot logical volume does not reside above the 4 GB address on the hard disk.

An example of this calculation follows:

1. Run **lsvg rootvg** to determine PP SIZE. On a 4.5 GB hard disk, the default PP SIZE is 8 MB. Make a note of that size.
2. Run **lslv -m** *bootlv00*, where *bootlv00* is the name of the newly created boot logical volume.

   The numbers in the second, fourth, and sixth columns indicate the physical partitions that have been assigned to the boot logical volume. If the PP SIZE is 8 MB, the boot logical volume must not use any physical partition above 511 (512 * 8 = 4096, which is 4 GB). Similarly, if the PP SIZE is 16 MB, the boot image must not use any partition above 255, and if the PP SIZE is 4 MB, the boot image must not use any partition above 1023.

Machine limitations with Universal Disk Format (UDF)

When booting a 7043-150 or 7046-B50 system from the Universal Disk Format (UDF) media, use the **O/F** command instead of SMS. The following is an example of how to use the **O/F** command:

```
boot /pci@fef00000/scsi@c/sd@4,0:1,\ppc\bootinfo.txt
```

Logical Volume Manager memory impact

The Logical Volume Manager (LVM) policies for allocating memory buffers were enhanced on AIX Version 5.3. These new memory allocation policies provide better LVM performance on AIX Version 5.3 and eliminate much of the need for tuning LVM parameters on AIX Version 5.3 versus previous versions of AIX. As a side effect of these enhancements, the pinned memory footprint of the LVM is larger on AIX 5.3 than on AIX 5.2 or AIX 5.1. Specifically, the LVM on AIX 5.3 might require as much as 4.4 MB of additional pinned memory per physical volume than it did on previous releases of AIX. Thus, on a system with two physical volumes, the LVM would require approximately 8.8 MB more of pinned memory on AIX 5.3 than on AIX 5.2 or AIX 5.1.

Much of the additional LVM memory requirements on AIX 5.3 are due to an enhancement related to memory affinity. The LVM on AIX 5.3 takes advantage of a system's memory affinity properties to improve performance. However, the performance benefit of the memory affinity enhancements tends to be small, and in environments that are constrained by memory, the impact of the additional AIX 5.3 LVM memory footprint may outweigh any added benefit of the memory affinity enhancements. Thus, for those environments that are memory constrained, the memory footprint of the AIX 5.3 LVM can be significantly reduced by disabling memory affinity for a system via the memory_affinity vmo tunable. For example, the following command disables memory affinity on a system:

```
vmo -r -o memory_affinity=0
```

After running the above command, the **bosboot** command must be run, and a system must be rebooted in order for the tunable change to take effect. In order to reduce the AIX 5.3 LVM memory footprint by disabling memory affinity, AIX 5.3 Technology Level 5300-05 must be installed.

With AIX 5.3 Technology Level 5300-05, the extra LVM memory footprint related to memory affinity only applies to hardware platforms that support memory affinity. Thus, disabling the memory_affinity tunable on a system, which doesn't support memory affinity, will not have any effect on the LVM memory footprint. On

systems where memory affinity is unsupported, the LVM only requires approximately 270 KB of additional pinned memory per physical volume on AIX 5.3 Technology Level 5300-05 than on AIX 5.2 or AIX 5.1.

On systems that support memory affinity, disabling memory affinity reduces the additional AIX 5.3 LVM pinned memory requirements to approximately 270 KB per physical volume. Thus, on a system with two physical volumes and memory affinity disabled, the LVM only requires an additional 540 KB on AIX 5.3 than on AIX 5.2 or AIX 5.1.

Chapter 3. Installation, migration, upgrade, and configuration information

Installation

This section contains information about installing AIX 5.3 that supplements the information contained in the AIX 5.3 installation documentation.

The following publications describe AIX 5.3 installation:

- *AIX 5L Version 5.3 Operating System Installation: Getting Started* (SC23-4940-02)
- *Installation and migration* (SC23-4887-03)
- *AIX 5L Version 5.3 AIX Installation in a Partitioned Environment* (SC23-4926-02)

The installation guides are available online in the AIX Information Center and in printed hardcopy.

To order these installation guides, contact your point of sale, or in the U.S., call IBM Customer Publication Support at 1-800-879-2755. Give the order number of the book you want to order.

To obtain AIX 5L Version 5.3 installation hints and tips, visit the Subscription Service at the following Web site:

`https://techsupport.services.ibm.com/server/pseries.subscriptionSvcs`

Installing AIX 5L Version 5.3

The following methods can be used to install AIX 5L Version 5.3:

- Complete overwrite installation
- Preservation installation
- Migration installation

Note: After you install or migrate a system to AIX 5L Version 5.3, you can install a lower level of AIX by restoring a system backup or by performing a new and complete overwrite with base media. Preservation installations from AIX 5L Version 5.3 to a lower level of AIX are not supported.

If your system has AIX 5.3 with 5300-00 through 5300-04 installed (you can verify the level by running the **oslevel -r** command), then you can use the base media or the Update CD to update to AIX 5L with 5300-05. In either case, use the **smitty update_all** command to perform the update.

Note: Because only the base install images are on the media, if you use the product media to update to AIX 5L with 5300-05, you cannot reject the software and return to the previous level.

To install AIX 5L Version 5.3, boot your system from the product media, and follow the instructions in the *Installation and migration* in the AIX Information Center.

Note: AIX 5L Version 5.3 cannot be installed on MCA (Micro Channel® Architecture) or PowerPC Reference Platform® (PReP) machines.

Disk format requirement

You cannot install AIX on an improperly formatted SCSI disk. AIX requires the disk to be formatted to a sector size supported by the attached SCSI controller. All AIX SCSI controllers support 512 byte sector SCSI disks. The 522 byte sector SCSI disks are only supported when they are attached to SCSI RAID controllers. If the disk has been formatted for SCSI RAID, but is not attached to a SCSI RAID controller,

the disk might not configure. If the disk does configure, it might be unreadable in the AIX environment. In some instances, the certify function and the format function in AIX diagnostics can be used to reformat the disk for the attached SCSI controller.

Creating a bootable CD

It is recommended that you create a CD that can be used to boot and perform maintenance on your system that matches your current level of AIX.

To create a bootable CD, run the following commands (where *cdx* is an attached CD writer). The **bosinst.data** file must be set for a prompted install (`PROMPT = yes`).

```
cd /var/adm/ras
ls ./bosinst.data ./image.data | backup -ivqf/tmp/fakemksysb
mkcd -m /tmp/fakemksysb -d /dev/cdx
```

OR

```
 mkcd -m /tmp/fakemksysb -S
```

will create a CD image that can be transferred to a system with a CD writer.

Note: The final **mkcd** command in the previous example makes an image that can be transferred to another system (AIX or non-AIX) for burning.

Base Operating System installation options

The information in this section supplements the ″Installation Options″ chapter of the *Installation and migration* in the AIX Information Center.

In the Base Operating System installation menus, if there are more than 50 disks on the system, the disks are ordinarily grouped by adapter. However, for some types of disks, the grouping is slightly different:

SCSI disks
Disks may be grouped by adapter or SCSI bus

IBM TotalStorage® DS4000
Disks are grouped by disk array controller (DAC)

In each case, the user can select the adapter, SCSI bus, or DAC by name and see the associated disks. The physical location of the adapter, SCSI bus, or DAC is also displayed.

rsct.opt.storagerm fileset

The **rsct.opt.storagerm** fileset is not automatically installed with the Reliable Scalable Cluster Technology (RSCT) updates. You can install this fileset after the RSCT updates are applied. If you install the **rsct.opt.storagerm** fileset, to reject the RSCT updates, you must uninstall the **rsct.opt.storagerm** fileset before you request PTF rejects.

AIX Toolbox Media and NIM lpp_sources

When a Network Install Manager `lpp_source` is used for base system installs, do not copy the contents of the *AIX Toolbox for Linux Applications* CD into the `lpp_source`. This results in multiple copies of **cdrecord** and **mkisofs** software installing during base installation. Neither **cdrecord** or **mkisofs** installs if multiple copies are present.

Graphics software bundle requires two CDs

Due to space constraints on the AIX base product media CDs, the graphics software bundle is now included on the *Volume 2* CD. For installations using CD media, you are prompted for the *Volume 2* CD if you install with the defaults for an overwrite or preservation installation (`Graphics Software = yes`).

If you create a Network Installation Management (NIM) lpp_source, you will not be prompted for the *Volume 2* CD to add the graphics software to the lpp_source. To add the graphics software after creating the lpp_source:

1. Type the **smitty nim_update_add** command. The **Add Software to an lpp_source** menu is displayed.
2. Select the **Graphics** bundle for **INSTALLP BUNDLE containing packages to add**.

Note: When creating an lpp_source from a prior release of AIX, and copying the **Graphics.bnd** software to the lpp_source, restore the **Graphics.bnd** appropriate to the release you are creating the resources to, and make it into a NIM resource. You can get the bundle from the **bos** image in the lpp_source:

```
# cd /tmp # restore -xvqf <lpp_source_path>/bos ./usr/sys/inst.data/sys_bundles/Graphics.bnd
```

This command restores in **/tmp/usr/sys/inst.data/sys_bundles**. Copy it to the location of your choice, and create a new NIM bundle resource, smitty nim_mkres.

Network Installation Management

Network Installation Management (NIM) includes a readme file that is installed with the NIM Master **bos.sysmgt.nim.master** fileset. The path name of the file is **/usr/lpp/bos.sysmgt/nim/README**. The readme file contains additional information about the AIX 5L Version 5.3 NIM product and includes the following topics:

- Restrictions on SPOT Creation for Releases Prior to 5.3 (New LPP_SOURCE Directory structure)
- Web-based System Manager NIM May Have Problems Installing SW on Client Machines
- Restrictions on Customize Operation for RPM Packages
- Steps Necessary For Adding GNOME -or- KDE Desktop Support

CSM Server

Before you install the CSM Server, you must read the *Software Planning and Installation Guide* for specific procedural steps to use when you are installing CSM. Installation of this product is not possible if you do not use the documented procedures in this book.

CSM Server requires four open-source rpm filesets that must be installed prior to installing the CSM Server software and its dependent packages, the CSM Distributed Command Execution Manager (DCEM) GUI and the CSM DCEM Web-based System Manager application. The required rpm filesets are:

- tcl
- tk
- expect
- conserver

As the root user, do the following to install the rpm filesets and the CSM Server:

1. Install the above rpm filesets using SMIT Install Software (type **smitty install_latest** at the AIX command line).
   a. Press F4, and select **/dev/cd0** (CD Drive) as the **INPUT device / directory for software**.
   b. Press F4, and select (by pressing F7 for each package) **tcl-8.3.3**, **tk-8.3.3**, **expect-5.32**, and **conserver-7.2.4** as the **SOFTWARE to install** values, or press F4 and select all the software packages you want.
   c. To accept the software license agreements, press the Tab key to change no to yes as the **ACCEPT new license agreements** value, and then press Enter.
   d. Press Enter again to confirm that you want to continue the installation process.
   e. Review the installation results, and press F3 to return to the installation panel, or press F10 to return to the AIX command line.

2. Install the CSM Server and its dependent software using SMIT Install Software (type **smitty install_latest** at the AIX command line).
   a. Press F4, and select (by pressing F7 for each package) **csm.server**, **csm.gui.dcem**, and **csm.gui.websm** as the **SOFTWARE to install** values, or press F4 and select all the software packages you want.
   b. To accept the software license agreements, press the Tab key to change `no` to `yes` as the **ACCEPT new license agreements** value, and then press Enter.
   c. Press Enter again to confirm that you want to continue the installation process.
   d. Review the installation results, and press F3 to return to the installation panel, or press F10 to return to the AIX command line.

Other installation information

Installation packaging formats

AIX 5L Version 5.3 supports the following installation-packaging formats:

- installp, AIX system installation command and packaging format
- RPM, a Linux installation command and packaging format
- ISMP, InstallShield Multi-Platform packaging format

With the **geninstall** command, you can list and install packages from media that contains installation images packaged in any of the listed formats. The **geninstall** and **gencopy** commands recognize the non-installp installation formats and either call the appropriate installers or copy the images, respectively.

The AIX 5L Version 5.3 product media contains installp packages and RPM packages that are installed during a BOS installation. The installp packages are located in the following path, where *mount_point* is the mount point:

`/mount_point/installp/ppc`

The RPM packages are located in the following path, where *mount_point* is the mount point:

`/mount_point/RPMS/ppc`

If you have media that contains ISMP packages for AIX 5.3, the ISMP packages are located in the following path, where *mount_point* is the mount point:

`/mount_point/ismp/ppc`

The **installp**, **bffcreate**, **geninstall**, **gencopy** and **nim** commands recognize this media structure.

For more information about software packaging, see the Software Product Packaging Concepts section in the *Installation and migration* in the AIX Information Center.

Migration

Maximum size of boot image increased

For AIX 5L Version 5.3, the maximum size of the boot image has changed from the previous value of 11,984 KB (12 MB minus 16 KB) to 31,984 KB (32 MB minus 16 KB).

System migrated to AIX 5.3 might experience double boot

When booting AIX 5L Version 5.3 on a system that has previously been running an earlier release of AIX, you may notice that the system automatically reboots and restarts the boot process. This is how the firmware processes changed information in the boot image.

This reboot also occurs if the process is reversed. A system previously running AIX 5.3 that is booting a release of AIX prior to 5.3 goes through the same process. This ″double boot″ occurs only once; if the stored value does not change, then the second boot does not occur.

If you install AIX 5.3 and continue to use only AIX 5.3, this double boot occurs once, and it occurs only if your system was running a pre-AIX 5.3 release before you boot AIX 5.3. Systems that are preinstalled with AIX 5.3 and use only AIX 5.3 do not experience the ″double boot.″

Migration from AIX 4.2.1

A system running AIX 4.2.1 must be updated with the September 1999 or later Update CD before migrating to AIX 5.3. The CD label should have the number LCD4-0252-13 or higher. To check a running system, verify that the **bos.rte.install** file is at level 4.2.1.17 or higher.

Replacements for the vmtune and schedtune commands

When you migrate a system from a previous version of AIX to AIX 5L Version 5.2 or AIX 5L Version 5.3, it is automatically set to run in compatibility mode (pre520tune mode). Compatibility scripts that replace the **vmtune** and **schedtune** commands are included with AIX 5.2, which means that the previous behavior of the tuning commands is preserved to a large extent after a migration.

However, when you migrate to AIX 5.3, the pre520tune compatibility mode applies only to settings that were configured with the **no** and **nfso** commands because the **vmtune** and **schedtune** commands are no longer included. The compatibility mode is meant to be a temporary help in the migration to the new tuning framework and should normally not be used with releases after AIX 5.2.

For more information about migrating your settings to the new commands, see ″Replacements for the vmtune and schedtune commands″ in the AIX Information Center.

xlC.rte fileset

If you are migrating to AIX 5.3 from AIX 4.2.x or AIX 4.3.x, check the level of the **xlC.rte** fileset by typing the following command:

```
lslpp -L xlC.rte
```

If the **xlc.rte** level is earlier than 5.0.2.x, you must apply APAR IY17981 before migrating to AIX 5.3. Without APAR IY17981 installed, the migrated system might fail to boot.

APAR IY17981 is available from the following Web site:

`http://www.ibm.com/servers/eserver/support/pseries/aixfixes.html`

bos.clvm.enh fileset after migration to AIX 5L Version 5.3

The **bos.clvm.enh** fileset is not installed when your system is migrated to AIX 5.3. After the migration is complete, users of the **bos.clvm.enh** fileset will need to re-install the fileset from the AIX 5.3 installation media.

KDE desktop

If the CDE and KDE desktops are both installed on a system migrated from AIX 4.3 to AIX 5.3, the KDE desktop might not start from the CDE login. To fix this problem, remove the following CDE startup information from the **/etc/inittab** file:

```
dt:2:wait:/etc/rc.dt
```

Note: You must have root user authority to remove this CDE startup information.

Do not delete the following KDE entry from the **/etc/inittab** file:

```
kdm:2:once:/opt/freeware/kde/bin/kdm
```

Performance monitoring API

The Performance Monitoring API is contained in the **bos.pmapi** fileset. A beta version of the same code was made available to selected customers, and also through alphaWorks®, under the name *pmtoolkit*.

The **bos.pmapi** fileset does not support the RS64-I (A35) processor. If you try to install the fileset on a machine with this processor, the installation fails and returns the following error:

```
setup_branchtable: Processor not yet supported.
instal:  Failed while executing the ./bos.pmapi.pmsvcs.post_i script.
```

When you are migrating from any level of AIX with any level of the beta fileset installed, you must uninstall the **pmtoolkit** fileset and reboot the machine before you install the **bos.pmapi** fileset. If you do not, the machine will fail to boot when you attempt to load the **pmtoolkit** fileset's kernel extension.

Verify that the **pmtoolkit** fileset is installed by typing the following at the command line:

```
lslpp -l pmtoolkit
```

- If you get the following output:

  ```
  lslpp: 0504-132  Fileset pmtoolkit not installed
  ```

  you can safely install the **bos.pmapi** fileset.

- If you get the following output:

  ```
    Fileset                      Level  State      Description
    ----------------------------------------------------------------------------
  Path: /usr/lib/objrepos
    pmtoolkit                  1.3.1.6  COMMITTED  Performance Monitor Toolkit
                                                   1.3.1
  ```

  complete the following steps:

  1. Run the following command:

     ```
     installp -u pmtoolkit
     ```

  2. Reboot the machine. After the machine reboots, you can safely install the **bos.pmapi** fileset.

SNMPv3

After you migrate to AIX 5.3, the non-encrypted version of SNMPv3 will run by default. If you have your own community, trap, or smux entries in your **/etc/snmpd.conf** file, those must be manually migrated to the **/etc/snmpdv3.conf** file. For instructions on how to migrate this information, see ″Network Management″ in *Networks and communication management* in the AIX Information Center.

Kerberos

All of the secure remote commands use the Kerberos Version 5 library and the GSSAPI library provided by IBM Network Authentication Service Version 1.4 that is located on the *AIX 5L Version 5.3 Expansion Pack* CD. However, you must install the **krb5.client.rte** fileset.

If you are migrating to AIX 5.3 and have Kerberos Version 5 installed, the installation scripts will prompt you to install the **krb5.client.rte** fileset. The secure remote commands support Kerberos clients and servers from both Native Kerberos 5 and DCE.

For more information, see ″Understanding the Secure Rcmds″ in *Managing Communicatoins and Networks* in the AIX Information Center.

AIX Toolbox for Linux Application migration information

If you previously installed the AIX Toolbox for Linux Applications and the level of the **rpm.rte** fileset is lower than 3.0.5.20, remove that software from the system before migrating to AIX 5.3. The Toolbox

software installed with **rpm.rte** levels prior to 3.0.5.20 are incompatible with software from the AIX Toolbox installed on AIX 5L Version 5.3 because of shared library restructuring.

Remove the software if you are performing a preservation installation and you established an **/opt/freeware** file system for the Toolbox software. The files in that file system will not be automatically overwritten during a preservation installation. To remove your existing rpm filesets, use the **destroyRPMS** tool available in the **/contrib** directory on the *AIX Toolbox for Linux Applications* CD by typing the following:

```
mount -vcdrfs -oro /dev/cd0 /mnt
/mnt/contrib/destroyRPMS
```

If you are migrating your system from AIX 4.3.3 to AIX 5L and you installed the **rpm.rte** fileset without creating your own **/opt** or **/opt/freeware** file system, after running the **destroyRPMS** command it is recommended that you remove the **/opt/freeware** directory and the **/usr/opt/freeware** directory before migrating. On AIX 5L, the system provides a **/opt** file system into which the **rpm.rte** fileset is normally installed. However, if the RPM Package Manager (RPM) finds a pre-existing **/usr/opt/freeware** directory, it uses this location instead. You do not need to do this if you want your RPM freeware installed under the **/usr** file system, but the **/opt** file system is recommended.

If you have already migrated with the **/usr/opt/freeware** file system and want to change this afterwards, run the **destroyRPMS** command again, remove any existing **/usr/opt/freeware** and **/opt/freeware** directories, and install the **rpm.rte** fileset again.

Additional information is also available on the *AIX Toolbox for Linux Applications* CD in the **/README.TXT** file.

Multipath I/O (MPIO)

After you migrate to AIX 5.3, some disk devices will no longer be configured as *other FC disk*. These devices instead will be configured as *MPIO other FC disk*. The affected devices are EMC SYMMETRIX, HDS OPEN, and IBM TotalStorage disk subsystems. These devices are configured as MPIO devices if the device was previously configured as *other FC disk*.

The following describes some of the similarities and differences that are seen after the device has migrated to an *MPIO other FC disk*.

Terminology:

- A path is each physical connection between the host system and the device.
- A path control module (PCM) is a device specific module that manages a device's I/O across its paths.

A device configured as *other FC disk* has the following properties:

- Contains multiple device instances created for each path the device was detected on.
- Supports user-changeable device attributes.
- Can migrate to a vendor-specific device when vendor-supplied, device-specific ODM pre-definitions are installed.
- Is meant to be a transitory state during boot and install. The vendor-specific device ODM pre-definitions should be installed before using the device in a production environment.

A device configured as *MPIO other FC disk* has the following properties:

- Contains only one device instance created and multiple path instances created. Also contains one path instance for each physical connection between the host system and the device.
- Supports user-changeable device attributes. There may be additional attributes that are PCM specific.
- Can migrate to a vendor-specific device when vendor-supplied, device-specific ODM pre-definitions are installed.

- Presently is not supported by PowerPath, MDS, or SSD path management products. To support any of these products the vendor-specific no-MPIO ODM pre-definitions must be installed. Attempting to control a device configured as an MPIO device will produce undetermined results. Data integrity issues will exist if the device is operated in this configuration.
- Is supported in a production environment. Device-specific vendor ODM pre-definitions are not required to be installed before using in a production environment.
- Allows for installing and booting to an MPIO device.

Migration issues

The following describes migration issues if the *MPIO other FC disk* support is removed after devices have been configured as *MPIO other FC disk*:

Migrating to an *other FC disk* can occur if the *MPIO other FC* support is removed. In this case, where the update is uninstalled with the force option, the AIX 5.3 release of MPIO handles the migration. If the system is rebooted, the device instance is in the define state. During migration, the device instance is left in the define state, and a new *other FC disk* instance is created.

If the system is not rebooted and the device instance is in the define state, the device instance is left in the define state, and a new *other FC disk* instance is created.

If the system is not rebooted and the device instance is in the available state, the device instance is unchanged.

There might also be *other FC device* instances created. If the *MPIO other FC device* is not open, an *other FC device* instance is created for each path the device is detected on. If the *MPIO other FC device* is in the open state, no *other FC device* instances are created. This is because the *MPIO other FC device* has already issued a **SCIOSTART** command to the FC adapter for each of the paths. The FC adapter does not allow two devices with the same worldwide name and worldwide nodename to exist in its internal data structures.

If *other FC device* instances were created, sending I/O to the device while it is configured as both *MPIO other FC* and *other FC device* can cause indeterminate device behavior or data damage. Reboot the system to correct this condition. After the system is rebooted, the *MPIO other FC device* instance will be in the defined state and can be removed using the **odmdelete** command. The **rmdev** command will not remove the device due to the missing pre-definitions.

For more information about MPIO, see the following AIX publications:

AIX 5L Version 5.3 System Management Concepts: Operating System and Devices in the section titled *Multipath I/O*.

AIX 5L Version 5.3 System Management Guide: Operating System and Devices in the section titled *MPIO Devices*.

System V Printing Subsystem migration from AIX 4.3.3

The System V Printing Subsystem is an alternate printing subsystem in AIX. The installation of the **bos.svprint.*** filesets in a TCB environment requires that the **lp** user ID (UID:11) and **lp** group ID (GID:11) are present in the system. Otherwise, the installation of these filesets will fail.

To avoid this problem, create the **lp** user (UID:11) and **lp** group (GID:11) accounts on the AIX 4.3.3 system prior to the migration.

Chapter 4. Limitations and restrictions

Base Operating System (BOS)

Known problems with the ksh and ksh93 commands

With the **ksh** command, when multiple shells have the **noclobber** option set and they redirect output to the same file, there could be a race condition that can result in multiple shell processes writing to the file. The shell does not detect or prevent such race conditions.

During login shell startup, the following files are processed in the order specified:

1. **/etc/environment**
2. **/etc/profile**
3. **.profile**
4. **.env**

IBM Directory with Ja_JP locale

In the Japanese environment, it is strongly recommended that you use IBM Directory in the Ja_JP locale. In other Japanese locales, the Server Administration GUI does not work properly.

System management

Cluster Systems Management

You can access the Cluster Systems Management (CSM) documentation from the following Web site:

`http://www.ibm.com/servers/eserver/clusters/library`

Click the **AIX cluster software documentation** link, and then click the **Cluster Systems Management** link.

Software requirements

The CSM management server can be any supported System p, iSeries™, or xSeries® machine. If the CSM management server is running AIX, it must use a minimum of AIX 5L Version 5.3 or AIX 5L Version 5.2 with the 5200-04 Recommended Maintenance package. Other machines within the CSM cluster are referred to as managed nodes. Managed nodes can also be any supported System p, iSeries or xSeries machine. If the managed node is running AIX, it must use a minimum of AIX 5L Version 5.3, AIX 5L Version 5.2 with the 5200-04 Recommended Maintenance package and APARs as previously mentioned, or AIX 5L for POWER Version 5.1 with the 5100-07 Recommended Maintenance package. Please refer to the CSM documentation for more information on which machines and which versions of Linux are supported in a CSM environment.

Reliable Scalable Cluster Technology

The Reliable Scalable Cluster Technology (RSCT) Resource Monitoring and Control (RMC) application is part of RSCT. The RSCT includes a readme file that is installed with the **rsct.core.utils** fileset. The file is located in the **/usr/sbin/rsct/README/rsct.core.README** directory and contains additional information about the RMC application.

Restriction for Japanese locales

When the responses specified with the predefined **notifyevent** script are used in Japanese locales, alphanumeric (English) characters should be used for the condition name. If the condition name has non-alphanumeric characters in the mail header, it will be damaged. To work around this problem, you can modify the **notifyevent** script to not use the **$ERRM_COND_NAME** environment variable in the mail subject.

Service Resource Manager (ServiceRM)

ServiceRM is an RSCT resource manager that creates serviceable events for problems found by AIX Diagnostics. ServiceRM sends these events to the Service Focal Point on the Hardware Management Console (HMC).

Web-based System Manager

Remote client management

An HTTP Server must be installed and configured using one of the following configuration methods:

- Installing the IBM HTTP Server 2.0.47.1 on an AIX machine
- Installing any other HTTP Server on an AIX machine

This is necessary to support remote client management using Web-based System Manager. Proper configuration of an HTTP Server allows an AIX machine to serve the remote client download pages, Java™ Web Start, applet pages, and online extended helps.

When installing the IBM HTTP Server 2.0.47.1 on an AIX machine:

- Use the **wsm_remote** Software Bundle (**smitty install_bundle**) to install the IBM HTTP Server.
- Upon successful installation of the software, the bundle's post-installation processing script consolidates the steps needed to configure and initiate remote access and file serving capabilities for Web-based System Manager.
- This software installation bundle prompts you to insert the AIX Expansion Pack media to install the IBM HTTP Server.
- If you obtained the IBM HTTP Server, Version 2.0.47.1 from the following IBM HTTP Server product Web site:

  `http://www.ibm.com/software/webservers/httpservers/`

  then the **wsm_remote** Software Bundle allows you to install IHS from the hard disk by specifying the directory path name that contains your copy of the software installation images. To install IHS on AIX using the **wsm_remote** Software Bundle, manually complete the setup as follows:

  1. The installation directory path name must be in the format `./ismp/ppc/package_name`. For example, downloaded installation images can be copied to the **/usr/sys/inst.images/ismp/ppc/IHS2** directory. In this example, the installation source name is **/usr/sys/inst.images**, and the package name is **IHS2**.
  2. The response file named **silent.res** must be linked to the name **IHS2.response** for AIX to detect automatic responses during a silent installation, such as specifying the **-P ihs.installLocation=/usr/HTTPServer** preferred AIX installation location and a language other than the default `en` (English) language.

When installing any other HTTP Server on an AIX machine, complete the following:

1. Install the Web server.
2. Upon successful installation of the software, configure the Web server using the **smitty web_based_system_manager** SMIT fast path command.
3. Provide the required information in the panels. A configuration script runs and consolidates the steps needed to configure and initiate remote access and file-serving capabilities for Web-based System Manager.

After updating the **sysmgt.websm** filesets, **wsmserver** in the **/etc/services** file and in the **/etc/inittab** file is updated. If modifications were previously made to these entries, the files might need to be re-edited after updating.

Inventory Scout, Version 2.2.0.9

Inventory Scout, Version 2.2.0.9 provides support for the new POWER5™ server family. The Vital Product Data (VPD) collection and formatting has changed significantly for this family of IBM servers, using the industry standard XML to encapsulate the VPD inventory data. These changes, for the most part, should be transparent to users of the system. The IBM tools and servers that receive VPD data have been enhanced to use this new format. The new XML-formatted VPD does not support the concatenation of VPD files that the previous format permitted.

Inventory Scout has a new microcode management graphical user interface (GUI). This feature is available on your AIX system by installing the **invscout.websm** fileset, or if a Hardware Management Console (HMC) is attached, by using the microcode update function. The GUI is a Web-based System Manager plug-in that can survey, download, and install the microcode levels of the system.

This release of Inventory Scout significantly changes the method used to determine the microcode levels of systems, adapters, and devices to compare it to the latest available levels. Previously, data was collected and sent to IBM to determine the state of the system.

The new microcode management feature does the following:

- Downloads a catalog of available levels to the system being examined
- Conducts a microcode survey on the system and compares it to the latest available microcode
- Allows you to download and flash to the latest microcode available for POWER4™ and POWER5 systems

The new microcode survey procedure might cause some problems with customer techniques used for surveying systems and might require changes to those procedures.

The microcode management feature relies on system features that were not present in previous generations of the systems. Support for microcode on these systems is limited to survey only. For more information about microcode updates, see the following Web site:

`http://www14.software.ibm.com/webapp/set2/firmware/gjsn`

To enable the new Inventory Scout functionality, the following filesets must be installed at the following levels or higher:

invscout.com	2.2.0.1
invscout.ldb	2.2.0.2
invscout.rte	2.2.0.9
invscout.websm	2.2.0.5

To obtain the required filesets, order APAR IY58377 from the following Web site:

`http://www.ibm.com/servers/eserver/support/unixservers/aixfixes.html`

Other software

This section contains information about other software. Additional information about AIX-supported products is available from the following Web sites:

- IBM Global Services Supported Products List (http://www.ibm.com/services/sl/products)
- IBM Software Support Lifecycle (http://www.ibm.com/software/info/supportlifecycle/)

Compilers

The following programs are fully supported versions:

- **C/C++**:

VisualAge C++ Professional for AIX, Version 6.0
C for AIX, Version 6.0
XL C/C++ Enterprise Edition, Version 7.0 for AIX
XL C Enterprise Edition, Version 7.0 for AIX
XL C/C++ Enterprise Edition, Version 8.0 for AIX
XL C Enterprise Edition, Version 8.0 for AIX

Information for current required APARs is available from the following web site:

`http://www.ibm.com/support/docview.wss?uid=swg21207318`

- **Fortran**:

  XL Fortran for AIX, Version 8.1.1
  XL Fortran Run-Time Environment for AIX, Version 8.1.1
  XL Fortran Enterprise Edition, Version 9.1 for AIX
  XL Fortran Enterprise Edition, Version 10.1 for AIX

  Information for current required APARs is available from the following web site:

  `http://www.ibm.com/support/docview.wss?uid=swg21207319`

- **COBOL**:

  COBOL for AIX, Version 2.0

For more information, see "Traditional Programming Languages" at the following web site:

`http://www.ibm.com/software/sw-bycategory/subcategory/SW760.html`

AIXlink/X.25

AIXlink/X.25, Version 2.1

AIXlink/X.25 Version 2.1 is supported on AIX 5.3.

For more information about supported adapters and about configuration and installation, see the *AIXlink/X.25 Version 2.1 for AIX: Guide and Reference* at http://publib.boulder.ibm.com/infocenter/pseries/index.jsp. Select **AIX documentation** → **Networking and communication** → **AIXlink/X.25 Version 2.1 for AIX: Guide and Reference**.

AIXlink/X.25, Version 1.1.5

AIXlink/X.25 Version 1.1.5 is not supported on AIX 5L Version 5.3.

AIX Fast Connect, Version 3.2

AIX Fast Connect documentation is available at the following Web site: http://publib.boulder.ibm.com/infocenter/pseries/v5r3/index.jsp?topic=/com.ibm.aix.doc/aixbman/fastcon/fastcon.htm. The latest updates of this product are described in the **/etc/cifs/README** file, which is installed with AIX Fast Connect.

Quick start

To install AIX Fast Connect:

1. Install AIX Fast Connect from the distribution CD using the **smitty install_all** fast path.
2. Use the **smitty smb** fast path to access AIX Fast Connect SMIT menus.
3. Configure AIX Fast Connect for encrypted passwords, and add a user.
4. Access the AIX Fast Connect server from a PC client by mapping a network drive. The server name is the same as the AIX host name, and HOME share is available by default.

Communications Server for AIX, Version 6.1

Communications Server for AIX, Version 6.1 requires PTF level 6.1.0.5 or later to run on AIX 5.3. See the Communications Server Support Web site to access the Communications Server for AIX, Version 6.1 PTFs:

`http://www.ibm.com/software/network/commserver/aix/support/`

You can download the PTF level 6.1.0.5 or later installable image from the Communications Server Service Update Web site using the Service Key provided with the product.

Note: The AnyNet® functions of CS/AIX are not supported on the 64-bit kernel.

Distributed Computing Environment (DCE)

IBM DCE for AIX, Version 3.2, support for AIX Version 5.3 requires PTF 7 or later.

IBM DFS™ for AIX, Version 3.1, support for AIX Version 5.3 requires PTF 9.

See http://www.ibm.com/software/network/dce/support/version/info.html for more details.

Enterprise Identity Mapping (EIM)

AIX 5.3 supports Enterprise Identity Mapping (EIM), which may be used to manage user identities across multiple systems. In some cases, EIM requires installation of Kerberos modules from the AIX 5.3 Expansion Pack. It is recommended that Kerberos be installed from the AIX 5L Version 5.3 Expansion Pack before using EIM.

Chapter 5. Documentation

AIX Information Center

The IBM System p and AIX Information Center is an information portal for AIX and System p customers. From this site, you can access the following:

- AIX 5L Version 5.1 for POWER, and AIX 5L Version 5.2 documentation
- Hardware documentation
- Message database for 7-digit error codes, LEDs, and error identifiers
- How-to's for users and system administrators
- FAQs
- Links to Redbooks™, white papers, and related products

To access the Information Center, go to the following Web site:

`http://publib16.boulder.ibm.com/pseries/index.htm`

Appendix A. AIX 5L Version 5.3 changes

New and Enhanced Functionality in the 5300-05 Technology Level

While not a major release, AIX 5L Version 5.3 with the 5300-05 Technology Level includes new functionality for security, reliability, systems management, performance and distributed file systems.

AIX 5L Version 5.3 now includes a new network hardening tool (AIX Security Expert) to help administrators increase the network security of their servers. Customers can also now use Microsoft® Active Directory as an LDAP directory and authentication server for AIX 5L systems.

AIX 5L Version 5.3 now virtualizes the Real Mode Area to provide for a more flexible and efficient use of real memory by partitions.

AIX 5L Version 5.3 extends the leadership capabilities of AIX 5L for reliability, availability and serviceability with new, mainframe inspired features such as:

- Faster systems dumps
- Networking, storage and **sysalloc** first failure data capture enhancements
- Kernel stack overrun protection and kernel no-execute page protection
- Trace enhancements for filesystems, networking, Virtual Memory Manager, Network Filesystem and other AIX 5L components
- Significant enhancements to **XMALLOC_debug** kernel service that can eliminate reboots to diagnose memory allocation problems

There are several systems management enhancements in this update:

- New commands to make it easier to boot multiple systems or partitions from a single NFS boot image which can greatly reduce the administrative workload associated with managing large, replicated AIX 5L environments.
- Better integration between SUMA and the Network Installation Manager (NIM) to gather and combine client inventories, consolidate download of AIX 5L fixes.
- Clients can also compare microcode, firmware, and AIX software inventories and generate reports.
- The **topas** tool has been enhanced to gather and record performance data from multiple partitions, simplifying performance management and capacity planning.
- AIX 5L now includes a new utility, the Network Data Administration Facility (NDAF). NDAF, based on IBM Research developed technologies, provides for easier administration of large NFS v4 distributed file systems.
- AIX 5L also includes new NFS proxy caching functionality that can greatly improve the efficiency of using NFS over Wide Area Networks without requiring changes to the NFS server or client.

Base Operating System (BOS)

64-bit system identifier

AIX 5.3 provides a 64-bit system identifier for compatibility with future systems.

64-bit kernel

AIX 5L Version 5.3 provides a scalable 64-bit kernel that is capable of supporting large application workloads running on 64-bit hardware. The 64-bit kernel scalability is primarily provided through a larger kernel address space. This space supports larger system software applications without requiring practical bounds and kernel extension interfaces.

32-bit kernel

The AIX 5L operating system previously contained both a uniprocessor 32-bit kernel and a 32-bit multiprocessor kernel. Effective with AIX 5L Version 5.3, the operating system supports only the multiprocessor kernel.

The AIX 5L Version 5.3 32-bit multiprocessor kernel supports the following systems: RS/6000, System p, or OEM hardware based on the Common Hardware Reference Platform (CHRP) architecture, regardless of the number of processors. The maximum real memory supported by a 32-bit kernel system (or partition) is 96 GB.

AIX 5L Version 5.2 is the last release of AIX that supports the uniprocessor 32-bit kernel.

Long user names, group names, and path names

User names, group names, and path names longer than eight characters are supported. Long name-enabled systems have interactive limitations with non-enabled systems.

The System Resource Controller (SRC Master) does not support the long user name, group name, and path name format in the following cases:

Client Application	SRC Master	Subsystem or Daemon
53	pre53	53
53	pre53	pre53
pre53	pre53	53

Notes:

1. The `53` notation is respective source code compiled with the AIX 5.3 OS environment.
2. The `pre53` notation is respective source code compiled with the pre-AIX 5.3 OS environment.

Any application calling **addssys()**, **chssys()**, **defssys()**, **getsubsvr()**, or **getssys()** should be recompiled with AIX 5.3 if it is to run on AIX 5.3.

System support

For information about supported and unsupported items, see Appendix B, “AIX 5L Version 5.3 unsupported devices,” on page 55.

Base functionality

The AIX 5L Version 5.3 kernels provide the same functionality, regardless of which kernel is being used. The 32-bit and 64-bit kernel systems have common base libraries, commands, utilities, and header files.

Differences between 32-bit and 64-bit kernel systems are limited to the following:

- **System and I/O Support.** The 64-bit kernel limits support to 64-bit POWER-based systems, while the 32-bit kernel supports both 32-bit and 64-bit POWER-based systems. In addition, the 64-bit kernel does not support all I/O that is supported by the 32-bit kernel.
- **Application Support.** The 64-bit kernel supports both 32-bit and 64-bit applications. Application source and binaries are portable between AIX 5L Version 5.3 64-bit and 32-bit kernel systems, in the absence of any application dependencies on internal kernel details or on kernel extensions that are not supported under the 64-bit kernel but are supported under the 32-bit kernel.
  - **Binary Compatibility.** Binary compatibility is provided for 32-bit applications running on earlier versions of AIX on POWER-based systems, except for applications linked statically or applications dependent on undocumented or unsupported interfaces. In addition, some system file formats have changed, and 32-bit applications processing these files might need to be recompiled.

- **Application Scalability.** AIX 5L Version 5.3 provides a more scalable application binary interface (ABI) for 64-bit applications. To take advantage of the scalability improvements to 64-bit programs, all 64-bit applications and libraries must be recompiled on AIX 5L Version 5.3. In addition, existing 32-bit kernel extensions and device drivers used by 64-bit applications might have to be modified in order to support the new 64-bit ABI.

- **Kernel Extensions.** Kernel extensions for the 64-bit kernel run in 64-bit mode and have the scalability of the larger kernel address space. Some kernel services available in the 32-bit kernel are no longer provided by the 64-bit kernel, so existing 32-bit kernel extensions may have to be ported in order to be used with the 64-bit kernel.

  Existing 32-bit kernel extensions continue to be supported by the 32-bit kernel, but these kernel extensions are not usable by the 64-bit kernel. Not all of the kernel extensions supported for the 32-bit kernel are supported for the 64-bit kernel, particularly the device drivers for the I/O.

  - **Dual-mode Kernel Extensions.** AIX 5L Version 5.3 supports dual-mode kernel extensions, which can be loaded by a common configuration method, regardless of which kernel is being used. A dual-mode kernel extension is an archive file that contains both the 64-bit and 32-bit versions of the kernel extension as members.

- **Installation and Enablement.** The 32-bit and 64-bit kernels are provided as part of the AIX 5L Version 5.3 base media and are installed on all supported hardware systems. The default kernel enabled during installation is dependent on the hardware system being installed. On POWER5 systems, the 64-bit kernel is enabled during base system installation. On all other systems, the 32-bit kernel is enabled. However, you can override this default option at installation time through the system installation panels.

  You can switch between the 32-bit and 64-bit kernels without reinstalling the operating system.

  1. Modify the **/usr/lib/boot/unix** directory and the **/unix** directory to be a symbolic link to the binary for the desired kernel.
  2. Run the **bosboot** command to write a new system boot image.
  3. Reboot the system.

  The path name of the 64-bit kernel is **/usr/lib/boot/unix_64**, and the path name of the multiprocessor versions of the 32-bit kernel is **/usr/lib/boot/unix_mp**.

JFS2 maximum file system support

JFS2 now supports file system sizes up to 32 terabytes (TB). The maximum file system size is dependent on the block size of the file system. The following table shows the maximum file system sizes for the various block sizes:

File System Block Size	Maximum File System Size
512	4 TB
1024	8 TB
2048	16 TB
4096	32 TB

JFS2 file system freeze and thaw feature

A new feature for the JFS2 file system is added to AIX 5L Version 5.3 with the 5300-01 Recommended Maintenance package. This feature provides an external interface whereby an application can request that a JFS2 file system freeze, or stay quiescent. After the freeze operation, the file system must remain quiescent until it is thawed or until the specified timeout has past.

The request for freeze or thaw can be performed from the command or from the API as follows:

- **Command:**

```
chfs -a freeze=<timeout or "off"> <file system name>
chfs -a refreeze=<timeout> <file system name>
```

- **API:**

```
fscntl()
fscntl(vfs, FSCNTL_FREEZE, (caddr_t)timeout, 0);
fscntl(vfs, FSCNTL_REFREEZE, (caddr_t)timeout, 0);
fscntl(vfs, FSCNTL_THAW, NULL, 0);
```

New memory allocation algorithm MALLOCTYPE=watson

AIX 5.3 introduces a new memory allocation algorithm, `MALLOCTYPE=watson`. The Watson malloc() setting can provide improvement over the default malloc in areas of memory fragmentation and performance in massively multithreaded applications, particularly with respect to small requests. New features are added to the malloc debugging facility to aid in the diagnosis of memory allocation problems, and malloc debugging capabilities are integrated into the dbx symbolic debugger.

For more information about the malloc debugging facility, see ″System Memory Allocation Using the malloc Subsystem″ in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs* in the AIX Information Center.

The -l option removed from the make command

The **make** command was enhanced to support "parallel execution" functionality. Two new flags **-j** and **-1** were added for this purpose. The permission/ownership of the **make** binary was changed to bin:system and the SGID bit was set. This was done to enable the **make** command to access the **/dev/kmem** file to get the load average for the **-l** option. This caused the LIBPATH to be erased for all SUID/SGID programs for security reasons.

The permission/ownership of the **make** binary has been restored so that LIBPATH is recognized. The **-l** option has been removed so it cannot be run by a user who is neither a root nor a member of system group.

The mksysb command

The method used by the **mksysb** command to restore data through system backups has changed.

Enhancements were added to more fully restore customized data so that a system more closely resembles the system at the time the backup was performed. This occurs when restoring a backup on the system that the backup originated from. These enhancements were added to reduce the amount of additional work that sometimes needs to occur to restore devices to their customized configuration at the time of backup.

If devices were removed from or replaced on the system after the backup was created, that information is restored when you are installing a backup, and the system shows these devices in a defined state.

These enhancements do not affect installing the backup onto other systems, or *cloning*.

Perl

Note: IBM continues to ship Perl, but it is an unsupported feature.

The following Perl filesets are included with AIX pursuant to the terms of the artistic license:

- **perl.rte** 5.8.2 (version 5.8)
- **perl.man.en_US**

For more information, run the **perl -v** command. To view the artistic license, see the following Web site:

`http://www.opensource.org/licenses/artistic-license.html`

The **perl.rte** fileset is automatically installed.

For more information about Perl, see the following Web site:
`http://www.perl.org`

The new Perl man pages are now located in the **/usr/share/man** directory.

AIX 5L Version 5.3 introduces Perl 5.8.2. If you have a Perl external subroutine compiled on earlier versions of Perl, the external subroutine may need to be recompiled with threading enabled on Perl 5.8.2.

C99 language interfaces

AIX 5L Version 5.3 system libraries and headers include interfaces required by the ISO/IEC 9899:1999(E) (C99) language standard and the Single UNIX Specification, Version 3. Some of the interfaces may have the same names as symbols in existing programs. The interfaces may be hidden by specifying the `-D_NOISOC99_SOURCE` when you are compiling.

Most of the new C99 language interfaces are unavailable when compiling to use the 128 bit long double floating point format rather than the default 64-bit long double format.

Domain errors generally do not occur for math routine error conditions.

Named shared library areas

By default, AIX shares libraries among processes using a global set of segments, referred to as the *global shared library area*. For 32-bit processes, this area consists of one segment for shared library text (segment 0xD) and one segment for pre-relocated library data (segment 0xF). Sharing text and pre-relocating data improves performance on systems where a large number of processes use common shared libraries.

Because the global shared library area is a single fixed-size resource, attempts to share a set of libraries that exceed the capacity of the area cannot succeed. In this situation, a portion of a process libraries are loaded privately. Loading libraries privately, as opposed to shared, consumes private address space in the process and places greater demands on paging space, leading to a degradation in overall system performance.

To address this limitation of the global shared library area, AIX 5.3 supports the following named shared library areas:

- A named shared library area replaces the global shared library area for a group of processes.
- A named shared library area enables a group of processes to have the full shared library capacity available to them at the same location in the effective address space as the global shared library area (segments 0xD and 0xF).
- The named shared library area feature is enabled via the **LDR_CNTRL** environment variable and no changes are required to existing binaries.
- Multiple named shared library areas can be active on the system simultaneously.
- Processes specify a particular named shared library area by a unique name. This name is chosen by the process that causes the area's creation.
- Named shared library areas are available for use only by 32-bit processes.

Because the use of a specific named shared library area is restricted to processes that request it, none of its space will be consumed by processes using the global shared library area or a different named shared library area.

Alternate memory model (doubletext32)

In addition to the default shared library area memory model (one segment dedicated to shared library text and one segment dedicated to pre-relocated library data), named shared library areas support an alternate memory model that dedicates both segments to shared library text. This model is useful for process

groups that share greater than 256 MB of library text. Note that since this alternate memory model performs no pre-relocation of library data, some performance degradation during module loading (for both exec-time dependencies and dynamically loaded modules) may be experienced. Therefore, the actual performance benefits of increased shared library text capacity should be considered on a case by case basis.

Interface

Access: A process requests the use of a named shared library area by having the **LDR_CNTRL** environment variable with the **NAMEDSHLIB** option in its environment at run time. The syntax of the new option is as follows:

`NAMEDSHLIB=`*`name`*`[,`*`attribute`*`][,`*`attribute2`*`]...[,`*`attributeN`*`]`

A valid *name* string can be any string matching the regular expression, `[A-Za-z0-9_\.]+` (containing only alphanumeric, underbar, and period characters).

A valid *name* string must be terminated by one of the following characters:

- `@` (at sign): The delimiter for multiple **LDR_CNTRL** options
- **`,`** (comma): The delimiter for **NAMEDSHLIB** attributes
- `\0` (null): The terminator of the **LDR_CNTRL** environment string

If an invalid *name* string is specified, the entire **NAMEDSHLIB** option is ignored. If an invalid *attribute* is specified, only that attribute is ignored. Currently, there is only one supported attribute: **doubletext32**.

There are no access restrictions for using named shared library areas. All requests for use of an area are granted.

Creation: There is no explicit interface to create a named shared library area. When a process requests the use of a named shared library area that does not exist, the area is automatically created.

Purging: The system automatically purges unused libraries from a named shared library area when it becomes full, just as it does for the global shared library area.

To force unused libraries to be purged from a named shared library area, a user runs the **slibclean** command with the **NAMEDSHLIB** option in the **LDR_CNTRL** environment variable. Root authority is required.

Destruction: There is no explicit interface to destroy a named shared library area. When the last process using a named shared library area exits (the usecount of the area drops to zero), the area is automatically destroyed.

Attributes: The **NAMEDSHLIB** attributes are examined by the system loader only during named shared library area creation. Therefore, requests to use an existing named shared library area are not strictly required to specify attributes matching those specified at creation (the request will not fail because of an attribute mismatch). However, because the system automatically destroys unused named shared library areas, it is good practice to always specify attributes, even when you are requesting the use of an existing named shared library area.

Examples

1. Run a pair of applications using the named shared library area named *XYZ* with one segment dedicated to shared library text and one segment dedicated to pre-relocated library data by running the following commands:

   ```
   $ export LDR_CNTRL=NAMEDSHLIB=XYZ
   $ xyz_app
   $ xyz_app2
   ```

2. Run a pair of applications using the named shared library area named *more_shtext* with both segments dedicated to shared library text by running the following commands:

   ```
   $ export LDR_CNTRL=NAMEDSHLIB=more_shtext,doubletext32
   $ mybigapp
   $ mybigapp2
   ```

3. Force a purge of the named shared library area named *XYZ* by running the following command:

   ```
   # LDR_CNTRL=NAMEDSHLIB=XYZ slibclean
   ```

4. Force a purge of the named shared library area named *more_shtext* by running the following command:

   ```
   # LDR_CNTRL=NAMEDSHLIB=more_shtext,doubletext32 slibclean
   ```

Geographic Logical Volume Manager (GLVM)

The Geographic Logical Volume Manager (GLVM) is software-based technology for realtime geographic data mirroring over standard TCP/IP networks. GLVM allows you to create a mirror copy of your data at a geographically distant location. Because of its tight integration with LVM, users who are already familiar with LVM should find GLVM very easy to learn. You configure geographically distant disks as remote physical volumes and then combine those remote physical volumes with local physical volumes to form geographically mirrored volume groups, that are managed by LVM in a manner similar to ordinary volume groups.

Documentation for GLVM is provided in two locations:

- The *HACMP/XD for Geographic LVM: Planning and Administration Guide* is available online at the following HACMP™ documentation page:

  `http://www.ibm.com/servers/eserver/pseries/library/hacmp_docs.html`

  This book provides complete planning, installation, configuration and usage information for GLVM in an HACMP/XD environment (GLVM was originally implemented with HACMP/XD for GLVM 5.2).
- Using standalone GLVM in AIX, apart from HACMP, is not covered in the above book. While many of the planning, configuration and usage steps are the same as for GLVM with HACMP/XD, there are a number of standalone procedures that are not covered in this book.

IBM 32-bit SDK for AIX, Java 2 Technology Edition, Version 1.4

IBM 32-bit SDK for AIX, Java 2 Technology Edition, Version 1.4 is released in **Java14.*** filesets. For more information, see the **/usr/java14/docs/sdkguide.aix32.htm** file.

IBM 32-bit SDK for AIX, Java 2 Technology Edition, Version 1.4 is included with the AIX base operating system. The 64-bit version is available on both the AIX 5L Version 5.3 Expansion Pack and at the following AIX Java Web site:

`http://www.ibm.com/developerworks/java/jdk/aix`

You can dynamically reconfigure a logical partition (LPAR) running a Java 1.4 application.

Note: Decreasing the number of processors or real memory allocated to an LPAR will likely degrade the performance of a Java application, but the application should continue to run.

IBM AIX Developer Kit, Java 2 Technology Edition, Version 1.3.1, 32-bit version for POWER and IBM AIX Developer Kit, Java 2 Technology Edition, Version 1.3.1, 64-bit version for POWER are both supported on AIX 5L Version 5.3. You can download these products from the AIX Java Web site. Install all of the Java service refreshes. To see if a more recent refresh is available, complete the following:

1. Go to the following developerWorks® Web site:

   `http://www.ibm.com/developerworks/java/jdk/aix`

2. Select the **Downloads, User Guides, and Service information** link.

3. In the table, select the **Fix Info** link from the **Java 1.3.1 32-bit** column or the **Java 1.3.1 64-bit** column.

As with Java 1.4, you can dynamically reconfigure an LPAR running Java 1.3.1.

The IBM SDK Version 1.3.1 supports user names of up to eight characters. Before running the **java** or **appletviewer** commands, ensure that you are logged in with a user name that is no longer than eight characters.

IBM 64-bit SDK for AIX, Java 2 Technology Edition, Version 5

IBM 64-bit SDK for AIX, Java 2 Technology Edition, Version 5 is released in **Java5_64.*** filesets. For the most current refresh available, perform the following steps:

1. Go to the developerWorks Web site at the following URL:
   `http://www.ibm.com/developerworks/java/jdk/aix`
2. Select the **Downloads, User Guides, and Service information** link.
3. In the table, select the **Fix Info** link from the **Java 5 64-bit** column.

System performance recordings and reports

AIX 5.3 now supports the following types of full-time recording of performance data:

1. **Local:** The **xmwlm** daemon (**-L** flag) records a limited set of system processor, memory, disk, network and other system metrics to generate a 24-hour log file.
2. **Multi-partition:** Data available from the **topas** cross-partition LPAR monitoring capability (**-C** flag) can now be recorded to a 24-hour log file using the **-R** flag.

   In this mode, the **topas** command operates as a background process and display functions are disabled.

For local recordings, data is placed in the **/etc/perf/daily/** directory and with files formatted as **xmwlm.YYMMDD**. The **/usr/lpp/perfagent/config_aixwle.sh** configuration script allows this function to be configured as an **inittab** process.

For multi-partition (LPAR) recordings, data is placed in the **/etc/perf/** directory and files are formatted as **topas_cec.YYMMDD**. The **/usr/lpp/perfagent/config_topas.sh** configuration script allows this function to be configured as an **inittab** process. If you are using the Performance Toolbox LPP, you must apply APAR IY76131 on any partition containing the **perfagent.server** fileset.

If you are using either type of recording, you must allocate additional space to the **/etc/perf** directory. As recordings are retained for 7 days, 1 MB/day should be allocated for each type of recording used (7 MB for each type).

For **topas -R** operation, choose a single LPAR to perform this data collection. The **topasout** command is used to post-process these recordings to generate text-based reports in formats similar to the **topas** command panels.

Configuration, use and detailed explanations of these functions are located in the **/usr/lpp/perfagent/README.perfagent.tools** file.

Reliability, availability, serviceability utilities

Reliability, Availability, Serviceability (RAS) is a collective term for those characteristics that enable a system to do the following:

- Perform its intended function during a certain period under given conditions
- Perform its function whenever it is needed
- Quickly determine the cause and the solution to a problem or error that affects system operation

The following sections highlight recent changes to some of the AIX RAS utilities and infrastructure.

System dump

Extended system failure status information is captured as part of the dump, detailing dump success or failure. Display this extended information by using the **sysdumpdev** command.

System dump compression is turned on by default. For information about dump compression, see the **sysdumpdev** command documentation.

In AIX 5L Version 5.3 with the 5300-05 Technology Level, there are changes to improve dump performance. As a result of these changes, the compression format of dump has changed. It is no longer a **.Z** compressed file. Instead it is a **.BZ** file (not to be confused with **bzip2** format which uses the **.bz2** extension). Use the **/usr/bin/dmpuncompress** command to uncompress the new dump format file. The **uncompress** command *does not* work on the new dump format file.

For example, for a compressed dump saved by snap or savecore, use:

```
dmpuncompress vmcore.1.BZ
```

If the dump is partial (non-zero status), the **-p** flag should be used to retrieve the partially compressed file:

```
dmpuncompress -p vmcore.1.BZ
```

For more information, refer to the **dmpuncompress** command documentation.

System dump is enhanced to support DVD-RAM as the dump media. A DVD-RAM can be used as the primary or secondary dump device.

The **snap** command is enhanced to support the following:

- Independent service vendors (ISVs) can use custom scripts to collect their custom problem data as part of the snap process. For programming and process details, see ″Copying a System Dump″ in *AIX 5L Version 5.3 Kernel Extensions and Device Support Programming Concepts*.
- Large outputs can be split into smaller files for ease of transport.
- Output can be written to DVD-RAM media.

In addition to any full system dump, a small minidump is now taken when the system crashes. The minidump is visible in the AIX error log after operating system reboot, and is included in the failure information sent to IBM service for diagnosis.

Advanced First Failure Data Capture features

AIX 5L Version 5.3 with the 5300-05 Technology Level package provides many advanced First Failure Data Capture (FFDC) features. These features include Lightweight Memory Trace (LMT), Component Trace (CT®), and Run-Time Error Checking (RTEC). These features are enabled by default, at levels that provide valuable FFDC information with minimal performance impacts. The advanced FFDC features can be individually manipulated. A System Management Information Tool (SMIT) dialog has also been provided, as a convenient way to persistently (across reboots) disable or enable the features through a single command.

To enable or disable all three advanced FFDC features, enter the following command:

```
smit ffdc
```

You can then choose to enable or disable FFDC features. Note that a **bosboot** and reboot are required to fully enable or disable all FFDC features. Any change to LMT will not take effect until the next boot.

System trace

The system trace facility has been enhanced to support process and thread-based tracing. You can restrict the tracing to a process and capture the events in relation to the process for better debugging. For more information, see the **trace** command documentation.

The **trace** command supports settings of larger trace buffers for regular users. For more information, see the **trcctl** command documentation.

The system trace can be used to trace processor utilization register (PURR) to provide more accurate event timings in a shared processor partition environment.

Lightweight Memory Trace

The Lightweight Memory Trace (LMT) provides system trace information for First Failure Data Capture (FFDC). It is a constant kernel trace mechanism that records software events occurring during system life. The system activates LMT at initialization, then tracing runs continuously. Recorded events are saved into per processor memory trace buffers. There are two memory trace buffers for each processor, one to record common events, and one to record rare events. The memory trace buffers can be extracted from system dumps and accessed on a live system by service personnel.

The impact on the throughput of a kernel-intensive benchmark is one percent, and is much less for typical user workloads. LMT requires the consumption of a small amount of pinned kernel memory. The default amount of memory required for the trace buffers is calculated based on factors that influence software trace record retention. For the 64-bit kernel, the default calculation is additionally limited such that no more than 1/128th of system memory can be used by LMT, and no more than 256 MB by a single processor. The 32-bit kernel uses the same default buffer memory size calculation, but restricts the total memory allocated for LMT (all processors combined) to 16 MB. The 64-bit kernel resizes the LMT trace buffers in response to dynamic reconfiguration events, the 32-bit kernel does not. The following table shows some examples of default LMT memory consumption:

Machine	Number of CPUs	System Memory	Total LMT Memory: 64-bit Kernel	Total LMT Memory: 32-bit Kernel
POWER3™ (375 MHz CPU)	1	1 GB	8 MB	8 MB
POWER3 (375 MHz CPU)	2	4 GB	16 MB	16 MB
POWER5 (1656 MHz CPU, SPLPAR, 60% ent cap, SMT)	8 logical	16 GB	120 MB	16 MB
POWER5 (1656 MHz CPU)	16	64 GB	512 MB	16 MB

To determine the amount of memory being used by LMT, enter the following shell command:

```
echo mtrc | kdb | grep mt_total_memory
```

The **raso** tunable command can be used to disable LMT. It can also be used to increase or decrease the memory trace buffer sizes. For more information, see the **raso** command documentation.

Component Trace

The Component Trace (CT) facility provides system trace information for specific system components. This information allows service personnel to access component state information through either in-memory trace buffers or through traditional AIX system trace. CT is enabled by default. The use of in-memory CT buffers can be persistently disabled across reboots by using the **ctctrl -P memtraceoff** command. CT can be persistently enabled by running the **ctctrl -P memtraceon** command.

Note: A **bosboot** is required to make the command persistent on the next boot

Information on these and other CT commands can be found in the **ctctrl** command documentation.

Run-Time Error Checking

The Run-Time Error Checking (RTEC) facility provides service personnel with a method to manipulate debug capabilities that are already built into product binaries. RTEC provides service personnel with powerful first failure data capture and second failure data capture error detection features. All Run-Time Error Checking can be persistently disabled across reboots by running the **errctrl -P errcheckoff** command. RTEC can be re-enabled persistently by running the **errctrl -P errcheckon** command.

Note: A **bosboot** is required to make the command persistent on the next boot.

For more information on the **errctrl** command, see *AIX 5L Version 5.3 Commands Reference, Volume 2*.

RTEC features include:

1. **Xmalloc debug**

   In AIX 5L Version 5.3 with 5300-05 Technology Level, random sampling of xmalloc allocations is enabled to catch memory leaks, buffer overruns and accesses to freed data. Xmalloc debug is similar to the previous memory overlay detection system (MODS). To specifically disable the xmalloc debug RTEC feature, run the **errctrl errcheckoff -c alloc.xmdbg -r** command. To enable xmalloc debug, run the **errctrl errcheckon -c alloc.xmdbg -r** command. For more information, see the MODS and **errctrl** command documentation.

2. **Excessive Interrupt Disablement Detection**

   The Excessive Interrupt Disablement Detection mechanism in AIX can detect whether or not privileged code remains disabled for interrupts for too long. Because excessive disablement might lead to performance problems, AIX writes an error log record to report this detection:

   ```
   IDENTIFIER TIMESTAMP  T C RESOURCE_NAME  DESCRIPTION
   A2205861   0705170705 P S SYSPROC        Excessive interrupt disablement time
   ```

   Report these error logs to IBM service. The detailed report contains additional information including a stack traceback and LMT (trace) data that can be used by IBM to identify the source of the problem.

   Only one period of interrupt disablement that exceeds .5 seconds is logged per boot (default). Note that each of these error log entries might identify a unique potential problem. These error reports are persistently disabled if RTEC is globally disabled. On a per-boot basis, disablement detection can be disabled by running the following command:

   ```
   errctrl errcheckoff -c proc.disa
   ```

   Finally, the following functions can be called from a disabled code section of a detected kernel extension to exempt the section from future excessive disablement reporting:

   ```
   disablement_checking_suspend
   disablement_checking_resume
   ```

   For more information about disablement checking, see ″disablement_checking_suspend Kernel Service″ and ″disablement_checking_resume Kernel Service″ in the *AIX 5L Version 5.3 Technical Reference: Kernel and Subsystems Volume 1*. Also see the **errctrl** command documentation.

Other RAS enhancements

The **chcore** command provides for management of location of core files. For more information, see the **chcore** command documentation.

AIX error logging now supports up to 4096 bytes of event data (see the **/usr/include/sys/err_rec.h** file). However, this size error log entry is intended only for restricted system use and general error log entries should continue to contain 2048 bytes or less of event data. While up to 4096 bytes of detail data is allowed, this size entry may be truncated across a reboot in certain circumstances. The largest detail data size guaranteed not to be truncated is 2048 bytes. A large error log entry reduces the non-volatile storage available to the system dump facility in the event of a system crash.

AIX Web browser transition to Mozilla

AIX 5L Version 5.3 supports the Mozilla Web Browser Version 1.7.0.12 (or later) as the default Web browser for AIX. A version of the browser is available to be ordered on CD media along with AIX. The latest version can be downloaded at no charge from the following Web site:

`http://www.ibm.com/servers/aix/browsers`

Mozilla for AIX requires GNOME libraries, which are available on the *AIX Toolbox for Linux Applications* CD or from the following Web site:

`http://www.ibm.com/servers/aix/products/aixos/linux`

Note: Netscape Communicator Version 4 is not supported on AIX 5.3.

Installing Mozilla for AIX

Mozilla for AIX can be installed as an option during the AIX Base Operating System installation process, or it can be installed later. All listed installation methods use the Mozilla installation bundle, which includes Mozilla and the required GNOME libraries.

The Mozilla installation process fails if the required GNOME libraries are not found. The required rpm filesets are listed.

Use one of the following installation methods:

- Install Mozilla using the following AIX BOS installation process:
  1. You can select Mozilla for installation during the AIX Base Operating System installation process by selecting these options in the following order:
     a. 2 = Change/Show Installation Settings and Install
     b. 3 = More Options
     c. 6 = Install More Software
     d. 1 = Mozilla (Mozilla CD)

     The default setting is to not install Mozilla.
  2. When prompted to do so, insert the *Mozilla* CD and the *AIX Toolbox for Linux Applications* CD.
- Install Mozilla as a bundle using the following Configuration Assistant process:
  1. Start **configassist**.
  2. Select **Manage software**, and click **Next**.
  3. Select **Install additional software**, and click **Next**.
  4. Select **Install by bundle**, and click **Next**.
  5. Specify the device or directory that contains the installation images, and click **Next**. If the location is a directory, such as **/usr/sys/inst.images**, verify the following:
     – The **Mozilla.base** installp package is in the **/usr/sys/inst.images/installp/ppc** directory
     – The toolbox rpm filesets are in the **/usr/sys/inst.images/RPMS/ppc** directory
  6. Select the Mozilla bundle, and click **Next**.
  7. Accept the license agreement, and click **Next** to start the installation process.
- Install Mozilla as a bundle using the following **smit** process:
  1. Run the **smit install_bundle** command.
  2. Specify the **INPUT device/directory** for software. If the location is a directory, such as **/usr/sys/inst.images**, verify the following:
     – The **Mozilla.base** installp package is located in the **/usr/sys/inst.images/installp/ppc** directory
     – The toolbox rpm filesets are located in the **/usr/sys/inst.images/RPMS/ppc** directory
  3. Select the **Fileset Bundle = Mozilla**.

4. In the Install Software Bundle screen, accept the license agreement, and press Enter to start the installation process.

Configuring Mozilla as the browser for AIX documentation services

Mozilla can be configured as the default browser that is used to view the AIX Documentation using Configuration Assistant or **smit**.

- Configure Mozilla using the following Configuration Assistant process:
  1. Start **configassist**.
  2. Select the Configure documentation server task.
  3. If Mozilla is detected as already installed, select **Yes, use Mozilla as the default browser**, and click **Next**.
- Configure Mozilla using the following **smit** process:
  1. Run the **smit change_documentation_services** command.
  2. Verify that **/usr/bin/mozilla** is set as the DEFAULT_BROWSER.

Migrating an existing Netscape Communicator Version 4 profile

If a Netscape Communicator Version 4 profile exists in your home directory and Mozilla is run for the first time, Mozilla prompts whether or not it should convert the Communicator profile, including the bookmarks to be used within Mozilla.

For more information about Mozilla for AIX, see the **/usr/mozilla/base/README.HTML** file.

License Use Management (LUM)

If your system has a 64-bit System ID (displayed by the **uname -f** command), the LUM Version 5.1 licensing software is installed. If your system has a nonzero 32-bit System ID (displayed by the **uname -u** command), the LUM Version 4 licensing software is installed. If both 32-bit and 64-bit System IDs are defined and the 32-bit ID is nonzero, both versions of LUM are installed.

LUM Version 4

The **i4blt**, **i4cfg**, **i4target**, and **i4tv** LUM Version 4 commands are in the **/usr/opt/ifor/ls/os/aix/bin** directory.

Note: To configure LUM Version 4, use the **i4cfg** command.

The default directory for the nodelock file is the **/var/ifor** directory.

For more information about LUM Version 4, see the *License Use Management User Guide* at **/usr/opt/ifor/ls/os/aix/doc/lumusg.htm**.

LUM Version 5

The **LUMblt**, **LUMcfg**, **LUMtarget**, and **LUMtv** LUM Version 5 commands are in the **/opt/LicenseUseManagement/bin** directory.

Note: To configure LUM Version 5, use the **LUMcfg** command.

The default nodelock directory is the **/var/LicenseUseManagement/nodelock** directory.

For more information about LUM Version 5, see the *License Use Management User Guide* at **/opt/LicenseUseManagement/doc/lumusg.htm**.

Exclusive resource sets

AIX Version 5.3 extends the support for resource sets to provide the option to create a set of exclusive use processors. The processors in an exclusive resource set (XRSET) are not globally available for use by any job on the system, and can only be used by jobs that explicitly attach to them using the attachment APIs.

If resource sets are being used in conjunction with Workload Manager (WLM), the number of processors in resource sets, exclusive or not, should be considered when defining processor shares and limits for classes. It is the responsibility of the administrator to ensure that the resource targets for each class are attainable within its set of resources.

Multiple instances of AIX on a single root volume group

In AIX 5.3, the root user can create multiple instances of AIX on a single root volume group (rootvg). A new utility, **/usr/sbin/multibos**, is supplied in AIX 5L with 5300-04 to create and manage a new instance of the operating system within the running rootvg. The **multibos** utility provides the root user operations to setup, access, maintain, update, and customize this new instance of the Base Operating System (BOS).

The result of creating a new instance of the BOS with **multibos** is a rootvg with two distinct and bootable instances of the operating system within a single rootvg. The running instance, called the active BOS, can be in production while **multibos** operations are used to modify the non-running instance, called the standby BOS.

The **multibos** command in the *AIX 5L Version 5.3 Commands Reference* incorrectly lists the supported level for **multibos** as 5300-02. You must run **multibos** with maintenance level 5300-04.

For more detailed information, refer to the latest **/usr/lpp/bos/README.multibos** file, and documentation regarding **multibos** in the AIX Information Center.

Multiple page size support

AIX 5L Version 5.3 with the 5300-05 Recommended Maintenance package includes 64-bit kernel support for virtual memory page sizes of 64 KB and 16 GB that are provided by POWER5+™ processors. These virtual memory page sizes are supported in addition to the previously supported virtual memory page sizes of 4 KB and 16 MB. Using a larger virtual memory page size like 64 KB for an application's memory can significantly improve an application's performance and throughput due to hardware efficiencies associated with larger page sizes.

The specific page sizes supported on a system depends on a system's processor type. You can use the **pagesize -af** command to display all of the virtual memory page sizes supported by AIX on a system.

You can specify the page sizes to use for three regions of a process's address space using an environment variable or settings in an application’s XCOFF binary with the **ldedit** or **ld** commands as shown in the following table:

Region	ld / ldedit option	LDR_CNTRL environment variable	Description
Data	-bdatapsize	DATAPSIZE	Initialized data, bss, heap
Stack	-bstackpsize	STACKPSIZE	Initial thread stack
Text	-btextpsize	TEXTPSIZE	Main executable text

For example, the following command causes **mpsize.out** to use 64 KB pages for its data, 4 KB pages for its text, and 64 KB pages for its stack on supported hardware:

```
$ LDR_CNTRL=DATAPSIZE=64K@TEXTPSIZE=4K@STACKPSIZE=64K mpsize.out
```

Unless page sizes are selected using one of the above mechanisms, a process will continue to use 4 KB pages for all three process memory regions by default.

Using 64 KB pages rather than 4 KB pages for a multi-threaded process's data can reduce the maximum number of threads a process can create. Applications that encounter this limit can reduce internal pthread library memory usage and allow for more threads to be created by setting the environment variable **AIXTHREAD_GUARDPAGES** to 0.

In addition to these three memory regions of a process's address space, you can select the page size for system V shared memory regions by using the **SHM_PAGESIZE** command to the **shmctl()** system call.

The 4 KB and 64 KB page sizes are intended to be general-purpose, and no system configuration changes are necessary to enable a system to use these page sizes. The 16 MB large page size and 16 GB huge page size are intended only to be used in very high performance environments, and a system administrator must configure a system to use these page sizes. Furthermore, the support for 16 MB large pages and 16 GB huge pages is limited. 16 MB large pages are only supported for process data and shared memory, and 16 GB huge pages are only supported for shared memory.

The **ps -Z** command displays the page sizes being used for the data, stack, and text memory regions of a running process. The **vmstat** command is enhanced to display information about multiple page sizes. The **-p** and **-P** options to the **vmstat** command displays VMM statistics for each supported page size.

Finally, the following **vmo** command can be used to disable all kernel support for 64 KB and 16 GB pages:

```
vmo -r -o vmm_mpsize_support=0
```

Communications, networking, and I/O

IP Security

AIX IP Security intrusion prevention system supports stateful filtering with a rich set of IF, ELSE, and ENDIF rules. It also guards against port scan-based attacks with a robust set of shun filters. Intrusion prevention is further strengthened with the ability to match and prevent patterns within the network data packets.

Using Certificate Management System (CMS) with Java 1.4

To use CMS, the following changes need to be made to **java.security** file, located in the **/usr/java14/jre/lib/security/** directory, when Java 1.4 is installed. CMS is part of the **AIX Certificate and SSL Base Runtime (GSKIT)** fileset that is included on the AIX 5L Version 5.3 Expansion Pack.

Locate the following stanza in the **java.security** file:

```
security.provider.1=com.ibm.jsse.IBMJSSEProvider
security.provider.2=com.ibm.crypto.provider.IBMJCE
security.provider.3=com.ibm.security.jgss.IBMJGSSProvider
security.provider.4=com.ibm.security.cert.IBMCertPath
```

Add the following two lines to the beginning of this stanza:

```
security.provider.1=sun.security.provider.Sun
security.provider.2=com.ibm.spi.IBMCMSProvider
```

The resulting stanza should show:

```
security.provider.1=sun.security.provider.Sun
security.provider.2=com.ibm.spi.IBMCMSProvider
security.provider.3=com.ibm.jsse.IBMJSSEProvider
security.provider.4=com.ibm.crypto.provider.IBMJCE
security.provider.5=com.ibm.security.jgss.IBMJGSSProvider
security.provider.6=com.ibm.security.cert.IBMCertPath
```

Write quit to save and close the file.

Then move the **gskikm.jar** file from the **/usr/java14/jre/lib/ext** directory to the **/tmp** directory.

Asynchronous I/O fast path for CIO with JFS2

Asynchronous I/O (AIO) allows applications to overlap processing and I/O operations to improve utilization of CPU and I/O resources. Concurrent I/O (CIO) allows multiple threads within an application to write concurrently to a single file while avoiding single-writer lock (inode lock) contention. Certain applications, especially transaction-oriented database applications, can take advantage of AIO and CIO to improve I/O throughput performance.

The AIO fast path for CIO is yet another I/O optimization. It allows I/O requests to be submitted directly to the disk driver strategy routine through the Logical Volume Manager (LVM). The fast path for CIO is supported exclusively with the JFS2 file system. Without the fast path, I/O must be queued to AIO kernel processes (kprocs) to handle the requests. In addition, the number of AIO kprocs must be tuned carefully to handle the I/O load generated by the application. The kproc path can result in slower performance than the fast path due to additional CPU/memory usage and inadequate AIO kproc tuning.

AIX Network Data Administration Facility

The AIX Network Data Administration Facility (AIX NDAF) provides for secure centralized management of file system data relationships across multiple systems. Its purpose is to facilitate central control of a federated file system namespace and replicated read-only data across a collection of AIX NFS Version 4 file servers. AIX NDAF is provided on the AIX 5.3 Expansion Pack media.

Internet Key-Exchange logging

The *Security* refers to the **/etc/isakmdp.conf** file instead of its correct name, **/etc/isakmpd.conf**. Use the **/etc/isakmpd.conf** file to configure these items:

- The log configuration of the IKE daemons.
- The manner in which the **isakmpd** command processes a proposed negotiation (whether the **isakmpd** daemon can accept a main-mode negotiation from an unknown peer).
- The manner in which the certificate revocation list (CRL) processes the following data:
  - The SOCKS4 server information
  - The LDAP server information
  - Whether the Hypertext Transfer Protocol (HTTP) server or the LDAP server is queried first, when both servers are configured

For more information, see ″Securing the Network → Internet Protocol security → Internet Protocol security problem diagnosis → Troubleshooting Key Exchange tunnel errors → Internet Key-Exchange logging″ in the *AIX 5L Version 5.3 Security Guide*.

RADIUS Server

The RADIUS Server implements a client and server protocol that lets remote access clients communicate with a central server to gain access to a network. The RADIUS server authenticates users, authorizes their requests for access to services, and writes accounting data. The protocol is based on IETF RFCs 2865 and 2866.

For more information, see ″RADIUS Server″ in the *Security*.

Path MTU (PMTU) discovery

Path Maximum Transmission Unit (PMTU)-related information is now stored separately from the routing table, in a table called the PMTU table. Routes are no longer cloned for IPv4, and the **netstat -rn** command no longer displays PMTU values. A new **pmtu** command is provided to view the PMTU table.

This command is used to display IPv4 and IPv6 entries and can also be used to delete a PMTU entry. A PMTU entry is added when the **route add** command runs with the `mtu` value specified. When a route is deleted, all PMTU entries using that route are also deleted. A **pmtu_expire** network option is provided to expire unused PMTU entries. The default value is 10 minutes.

User Datagram Protocol (UDP) applications using PMTU discovery must always specify the **IP_DONTFRAG** socket option, along with the **IP_FINDPMTU** socket option.

AF_INET6 sockets

Beginning with AIX 5.3, the behavior of **AF_INET6** sockets for protocol **IPPROTO_RAW** has changed to comply with RFC3542.

When an application performs a `receive` on this type of socket, it will receive *only* payload data from the packet. In earlier versions of AIX, when an application performed a `receive` on an **AF_INET6 IPPROTO_RAW** socket, it received the IPv6 header, followed by the payload data.

To preserve the former behavior (to continue receiving the IPv6 header followed by payload data), applications must now have their code modified to set the new **IPV6_AIXRAWSOCKET** socket option on any **AF_INET6 IPPROTO_RAW** sockets and recompile.

Mismatch in htonl function prototype

The function prototype of the **htonl**, **ntohl**, **htons**, and **ntohs** subroutines as given in manpages differs from the function prototype found in the **net/nh.h** file. The new **htonll** and **ntohll** subroutines are also listed.

htonl subroutine

The **htonl** subroutine converts a 32-bit unsigned integer from host byte order to Internet network byte order. The Internet network requires addresses and ports to be in network standard byte order. Use the **htonl** subroutine to convert the host integer representation of addresses and ports to Internet network byte order. The **htonl** subroutine is defined in the **net/nh.h** file as a null macro if the host byte order is same as the network byte order. The **htonl** subroutine is declared in the **net/nh.h** file as a function if the host byte order is not same as the network byte order.

All applications containing the **htonl** subroutine must be compiled with **_BSD** set to a specific value (acceptable values are 43 and 44). All socket applications must include the BSD **libbsd.a** library. The syntax is:

```
#include <sys/types.h>
#include <netinet/in.h>

uint32_t htonl ( HostLong)
uint32_t HostLong;
```

where *HostLong* is a 32-bit integer in host byte order. The **htonl** subroutine returns a 32-bit integer in Internet network byte order (most significant byte first).

htons subroutine

The **htons** subroutine converts a 16 bit unsigned integer from host byte order to Internet network byte order. The Internet network requires ports and addresses to be in network standard byte order. Use the **htons** subroutine to convert addresses and ports from their host integer representation to network standard byte order. The **htons** subroutine is defined in the **net/nh.h** file as a null macro if the host byte order is same as the network byte order. The **htons** subroutine is declared in the **net/nh.h** file as a function if the host byte order is not same as the network byte order.

All applications containing the **htons** subroutine must be compiled with **_BSD** set to a specific value (acceptable values are 43 and 44). In addition, all socket applications must include the BSD **libbsd.a** library. The syntax is:

```
#include <sys/types.h>
#include <netinet/in.h>

uint16_t htons ( HostShort)
uint16_t HostShort;
```

where *HostShort* specifies a 16 bit integer in host byte order that is a host address or port. The **htons** subroutine returns a 16 bit integer in Internet network byte order (most significant byte first).

ntohl subroutine

The **ntohl** subroutine converts a 32-bit unsigned integer from Internet network byte order to host byte order. Receiving hosts require addresses and ports to be in host byte order. Use the **ntohl** subroutine to convert Internet addresses and ports to the host integer representation. The **ntohl** subroutine is defined in the **net/nh.h** file as a null macro if the host byte order is same as the network byte order. The **ntohl** subroutine is declared in the **net/nh.h** file as a function if the host byte order is not same as the network byte order.

All applications containing the **ntohl** subroutine must be compiled with **_BSD** set to a specific value (acceptable values are 43 and 44). In addition, all socket applications must include the BSD **libbsd.a** library. The syntax is:

```
#include <sys/types.h>
#include <netinet/in.h>

uint32_t ntohl ( NetLong)
uint32_t NetLong;
```

where *NetLong* requires a 32-bit integer in network byte order. The **ntohl** subroutine returns a 32-bit integer in host byte order.

ntohs subroutine

The **ntohs** subroutine converts a 16 bit unsigned integer from Internet network byte order to host byte order. Receiving hosts require Internet addresses and ports to be in host byte order. Use the **ntohs** subroutine to convert Internet addresses and ports to the host integer representation. The **ntohs** subroutine is defined in the **net/nh.h** file as a null macro if the host byte order is same as the network byte order. The **ntohs** subroutine is declared in the **net/nh.h** file as a function if the host byte order is not same as the network byte order.

All applications containing the **ntohs** subroutine must be compiled with **_BSD** set to a specific value (acceptable values are 43 and 44). In addition, all socket applications must include the BSD **libbsd.a** library. The syntax is:

```
#include <sys/types.h>
#include <netinet/in.h>

uint16_t ntohs ( NetShort)
uint16_t NetShort;
```

where *NetShort* requires a 16 bit integer in network standard byte order. The **ntohs** subroutine returns a 16 bit integer in host byte order.

htonll subroutine

The **htonll** subroutine converts a 64-bit unsigned integer from host byte order to Internet network byte order. The Internet network requires addresses and ports to be in network standard byte order. Use the **htonll** subroutine to convert the host integer representation of addresses and ports to Internet network byte order. The **htonll** subroutine is defined in the **net/nh.h** file as a null macro if the host byte order is same as the network byte order. The **htonll** subroutine is declared in the **net/nh.h** file as a function if the host byte order is not same as the network byte order.

All applications containing the **htonll** subroutine must be compiled with _**BSD** set to a specific value (acceptable values are 43 and 44). In addition, all socket applications must include the BSD **libbsd.a** library. The syntax is:

```
#include <sys/types.h>
#include <netinet/in.h>

uint64_t htonll ( HostLong)
uint64_t HostLong;
```

where *HostLong* specifies a 64-bit integer in host byte order. The **htonll** subroutine returns a 64-bit integer in Internet network byte order (most significant byte first).

ntohll subroutine

The **ntohll** subroutine converts a 64-bit unsigned integer from Internet network byte order to host byte order. Receiving hosts require addresses and ports to be in host byte order. Use the **ntohll** subroutine to convert Internet addresses and ports to the host integer representation. The **ntohll** subroutine is defined in the **net/nh.h** file as a null macro if the host byte order is same as the network byte order. The **ntohll** subroutine is declared in the **net/nh.h** file as a function if the host byte order is not same as the network byte order.

All applications containing the **ntohll** subroutine must be compiled with _**BSD** set to a specific value (acceptable values are 43 and 44). In addition, all socket applications must include the BSD **libbsd.a** library. The syntax is:

```
#include <sys/types.h>
#include <netinet/in.h>

uint64_t ntohll ( NetLong)
uint64_t NetLong;
```

where *NetLong* requires a 64-bit integer in network byte order. The **ntohll** subroutine returns a 64-bit integer in host byte order.

Removal of support for devices

The following devices are not supported on AIX 5.3:

- **PCI FDDI I/O** (FC 2741, FC 2742, and FC 2743) is not supported on AIX 5.3.
- **devices.pci.b7105090**. The Ethernet adapter that is supported by the **devices.pci.b7105090** fileset in AIX versions *prior* to AIX 5L Version 5.1 is *not* supported in AIX 5L Version 5.3. After a migration to AIX 5L Version 5.3, or when AIX 5L Version 5.3 is installed and this Ethernet adapter is in the machine, the following messages may display on the console or be written to log files:

  ```
  Method error (/usr/lib/methods/cfgv3boom -l ent1 ):
          0514-068 Cause not known.

  cfgmgr: 0514-621 WARNING: The following device packages are required for
          device support but are not currently installed.

  devices.pci.b7105090 Not found on the installation media.
  ```

  Remove the unsupported Ethernet adapter from the machine. This adapter will not be configured by AIX 5L Version 5.3.

The devices.artic960 fileset

The **devices.artic960** fileset provides support for the following IBM ARTIC960 adapters:

- IBM ARTIC960Hx 4-Port Selectable PCI Adapter (FC 2947)
- IBM ARTIC960RxD Quad Digital Trunk Adapter (FC 6310)

This includes EEH support and 64-bit support for FC 2947 and FC 6310 adapters. If an additional fileset is installed to access a particular IBM ARTIC960 adapter, full EEH and 64-bit support depends on the ability of the additional fileset to support EEH and 64-bit.

Included with the **devices.artic960** fileset are the following filesets:

- **devices.artic960.rte**, IBM ARTIC960 Runtime Support
- **devices.artic960.ucode**, IBM ARTIC960 Adapter Software
- **devices.artic960.diag**, IBM ARTIC960 Adapter Diagnostics

When a PCI I/O error occurs on an IBM PCI ARTIC960 adapter, the adapter slot becomes frozen and the IBM ARTIC960 adapter can be reset. Following an EEH error, the adapter software needs to be downloaded to the adapter again.

To determine if an EEH error occurred on an IBM ARTIC960 adapter, inspection of the error log is necessary. A temporary EEH error on an IBM ARTIC960 adapter is logged as a temporary EEH error followed by I/O errors specific to the IBM ARTIC960 adapter. Recovery from a temporary EEH error is accomplished by removing and making the IBM ARTIC960 device driver using the **rmdev** and **mkdev** command. This process loads the necessary adapter software onto the adapter.

If the error log shows a permanent EEH error, it is necessary to use the hot plug manager to remove and make the adapter again.

The devices.pci.14108c00 fileset

The **devices.pci.14108c00** fileset provides support for SDLC and bi-synchronous protocols on the IBM ARTIC960Hx 4-Port Selectable PCI Adapter (FC 2947). When combined with the installation of the **devices.artic960** fileset, Enhanced Error Handling (EEH) support is provided. Either 32-bit or 64-bit kernel mode is supported. 32-bit applications are supported.

Missing resource processing

In a partitioned environment, missing resource processing (through the **diag -a** command) is not performed for processors, memory, L2 Cache, integrated devices, or pluggable adapters that have been moved to another partition. This is done to aid configuration for resources that are moved from one partition to another partition, then moved back to the original partition.

To remove a device from the configuration, log in as the root user, and type `rmdev -dl` *`device`* at a command prompt, where *device* is the name of the device you want to remove.

For more information, view the service hints section within diagnostics when you are logged in as the root user or using the CE login. You can view the service information by doing the following:

1. At the command line, type `diag`.
2. When **Diagnostic Operating Instructions** is displayed, press Enter.
3. At the Function selection menu, select **Task Selection**.
4. At the Task Selection menu, select **Display Service Hints** and press Enter.

IBM Tivoli Directory Server (LDAP)

To access the latest IBM Tivoli Directory Server 6.0 product information, go to the following Web site:

`http://www-306.ibm.com/software/tivoli/products/directory-server/req-aix.html`

Installation and configuration

For information specific to IBM Tivoli® Directory Server installation and configuration, go to the following Web site:

`http://www.ibm.com/software/tivoli/products/directory-server/`

From the **IBM Tivoli Directory Server** category on this Web site, click **Technical Documentation** > **Version: 5.2**. Read the following documents:

- *Installation and Configuration*
- *Server Readme*
- *Client Readme*
- *Readme Addendum*

Before you run the **ldapxcfg** command, the following symbolic links must exist:

```
/usr/ldap/db2 -> /usr/opt/db2_08_01
/usr/ldap/lib/libdb2.a -> /usr/opt/db2_08_01/lib/libdb2.a
```

You can verify that these links exist by typing the following commands:

```
ls -l /usr/ldap/lib
ls -l /usr/ldap/db2
```

If these links are not present, create these links by typing the following commands:

```
ln -s -f /usr/opt/db2_08_01/lib/libdb2.a /usr/ldap/lib/libdb2.a
ln -s -f /usr/opt/db2_08_01 /usr/ldap/db2
```

You can install the Web Administration Tool on a system with or without the client or server.

If you are using DB2® 8.1, you must enable asynchronous I/O before you begin the configuration. To enable asynchronous I/O, type the following command at the command prompt:

```
smitty aio
```

There is a size underestimation with **ldap.server.com**. This can cause the installation to fail if not enough disk space is allocated.

Administration

After you install the **ldap.client** package, create the following link by typing the following:

```
ln -s -f /usr/ldap/lib/aix5/libldapiconv64.a /usr/lib/libldapiconv64.a
```

Dynamic Tracking and Fast I/O Failure of Fibre Channel devices

AIX supports Dynamic Tracking and Fast I/O Failure of Fibre Channel devices.

Dynamic Tracking allows the user to perform certain prescribed storage area network (SAN) changes that result in N_Port ID changes (such as cable movement at the switch ports or the creating of inter-switch links) without taking devices offline.

Fast Fail causes I/Os down a particular link to fail faster due to lost links between the switch and the storage device. This may be useful in a multipath environment where you want I/Os to fail over to another path relatively quickly.

Independent Software Vendors (ISVs) developing kernel extensions and/or applications that communicate with the AIX Fibre Channel Driver stack should refer to the ″Fibre Channel Protocol for SCSI and iSCSI Subsystem″ article in *AIX 5L Version 5.3 Kernel Extensions and Device Support Programming Concepts* in the AIX Information Center for changes necessary to support Dynamic Tracking.

Note: Pay special attention to the *Required FCP and iSCSI Adapter Device Driver ioctl Commands* and *Understanding the scsi_buf Structure* sections.

Internet Protocol (IP) over Fibre Channel

This information supplements the ″Internet Protocol (IP) over Fibre Channel″ section in the *Networks and communication management* in the AIX Information Center.

Note: IP over Fibre Channel is supported only within one subnet. Routing through a gateway is not supported. IP packets cannot be sent to a different subnet over Fibre Channel.

To set up the Fibre Channel networking capability between AIX and Thomson Grass Valley™ Media Servers (such as Profile PVS 1000), note that feature 6228 (IBM 2 Gigabit Fibre Channel Adapter for 64-bit PCI Bus) is currently the only Fibre Channel adapter supported for this configuration. In addition, the following steps are recommended to establish the communication between the two host systems:

1. Disable FARP (Fibre Channel Address Resolution Protocol) on AIX. For example, if the IP over Fibre Channel protocol driver instance is `fcnet0`, type the following command:

   ```
   chdev -l fcnet0 -a enable_farp=no
   ```

2. If the Profile Media Server is set up to disable ″checksum and re-transmit on error″ (this is the default option), type the following command on the AIX command line, assuming `fc0` is the interface for this purpose:

   ```
   ifconfig fc0 tcp_disable_cksum
   ```

   You can also use the AIX `ifconfig fc0` command to verify whether TCP checksum is disabled. To re-enable the TCP checksum on the IP over Fibre Channel interface, type the following command:

   ```
   ifconfig fc0 -tcp_disable_cksum
   ```

3. It is recommended that initial communication always be established from AIX with the **ping** command. Initiating the exchange (pinging) from the Profile Media Server, prior to any exchange over Fibre Channel from the AIX side, might result in a prolonged delay in the establishment of communication between the two systems.

Sendmail, Version 8.13.4

AIX has been updated with Sendmail Version 8.13.4 that supports SSL encryption based on OpenSSL. The OpenSSL package is available on the AIX Toolbox for Linux Applications CD, DVD or can be downloaded from the following Web site:

`http://www14.software.ibm.com/webapp/iwm/web/preLogin.do?source=aixtbx&S_PKG=dlaixww`

Generic Routing Encapsulation

AIX now supports Generic Routing Encapsulation (RFC 2784) tunneling protocol that can be used to redirect Web traffic to AIX servers from anywhere in the network.

AIX iSCSI software initiator

Beginning with AIX 5.2 with 5200-04, the iSCSI protocol driver is included as part of AIX Base Operating System. The iSCSI protocol allows the access of storage devices over gigabit Ethernet TCP/IP networks. The iSCSI support is in the filesets **devices.iscsi_sw.rte**, **devices.iscsi.disk.rte**, and **devices.common.IBM.iscsi.rte**. These filesets supersede the **iscsi_sw.rte** fileset that was previously included in the AIX Bonus Pack.

To use the iSCSI protocol driver, add the names of the iSCSI targets being accessed to the **/etc/iscsi/targets** file. For more information about configuring iSCSI, see the ″iSCSI Software Initiator″ section in the *AIX 5L Version 5.3 System Management Guide: Communications and Networks*. For more information about the **/etc/iscsi/targets** file, see the *AIX 5L Version 5.3 Files Reference*.

The AIX iSCSI protocol driver supports the 1 port and 2 port gigabit Ethernet adapters with optical or copper connections (FC 5700, FC 5701, FC 5706, and FC 5707). The iSCSI protocol driver is verified to work with the Cisco MDS 9000 IPS module as the iSCSI target, attaching to the IBM TotalStorage ESS F20 and IBM TotalStorage ESS 800 storage devices.

The current iSCSI protocol driver implements the draft-20 version of the IETF iSCSI standard, with the following limitations:

- During installation, the iSCSI driver creates a default initiator name. However, this generated iSCSI name might not comply with the format specified by the iSCSI String Profile document. You can use the iSCSI SMIT panels (under **smit iscsi**) to change the initiator name to comply with the standard or to match local iSCSI name conventions.
- The iSCSI protocol driver can connect to a maximum of 16 unique targets at one time. If fewer targets are in use, you can change the **Maximum Targets Allowed** field in the SMIT panel to reduce memory usage by the iSCSI driver.
- This implementation of iSCSI supports only one TCP/IP connection per iSCSI session.
- This implementation of iSCSI supports login redirection to numeric IP addresses only. Any received login redirection that specifies a host name instead of a numeric IP address is considered a login failure.

Configurable IP Multipath Routing

The Configurable IP Multipath Routing Feature provides functionality and flexibility when you are using the IP Multipath Routing (MPR). With MPR, you can configure multiple routes to a single destination (network or host routes), as long as the gateways are different. MPR provides a degree of fault tolerance and helps in load balancing across multiple paths.

Currently, you can configure multiple default routes through different gateways. When multiple routes are configured to the same destination network or host, these routes are used in a round-robin fashion.

MPR incorporates new policies into the route selection process in addition to the default round-robin policy (a special case of Weighted Round-Robin when the weights are `1` or not configured). With this feature, you can select and configure any of the following policies to be used with MPR:

- Weighted Round-Robin
- Random
- Weighted Random
- Lowest Utilization
- Hash-Based

Each of these policies works as follows:

Weighted Round-Robin (WRR)
: You can configure the multiple routes to a destination network to have different weights. These weights will determine the manner in which these routes are used. For example, if you configure three routes of a multipath routing set (it is implied that all three routes have the same destination network or host but go through different gateways) to have weights of `3`, `5`, and `2` respectively, and then configure the policy to be WRR, then the first route will be used three times (three different connections) before moving on to using the second route, which will be used five times, and then the third route, which will be used twice. Then it uses the first route in the same manner as before. Therefore, the round-robin now uses the weights to perform Weighted Round-Robin selection. If the weights are not configured (default is `1`), then regular round-robin is used.

Random (RND)
: As the name suggests, a route is chosen at random from the multipath routing set.

Weighted Random (WRND)
: With this policy, the configured weights of all routes are added, and a random number between `0` and the total weight is chosen. This random number is scaled down to a number between `0` and the number of routes in the MPR set, and chooses the route that corresponds to this scaled-down value.

Lowest Utilization (LUT)
With this policy, a route with the lowest reference count in the MPR set is chosen. The reference count is an indicator of the number of active connections using this route and therefore is used as an indicator of the use of these routes.

Hash-Based (HSH)
With this policy, a hash calculation is performed, based on the destination IP address, and a route is selected. You should not use this policy in the following cases:

- If the routes in an MPR set are all host routes.
- If most connections from the host that is being configured are to a specific destination IP address.

In both cases above, because this policy is based on the destination IP address, the hash-algorithm always chooses the same route.

These new policies can be configured globally on a per-system basis, or on a per-MPR set basis (each MPR set comprises a set of routes to a single destination network or host). The local setting takes precedence over the global setting when it is configured.

The global configuration is provided through a new network option that can be viewed through the **no** command. The new option is called **mpr_policy**, which can be set to any value from `1` to `5`, with each value corresponding to the numbers above for the policies. For more information about this option, see the **no** command documentation.

To configure the policies on a per-MPR set basis (each MPR set comprises a set of routes to a single destination network or host) the policy can either be set during route creation or by using the **route set** command after the route has been created. For more information about this command, see the **route** command documentation and the examples below.

The configuration information, such as the weight and policies currently used, can be viewed using the **netstat -Cn** command.

Examples

Example 1: Adding Multiple Routes with Different Weights: To add multiple routes with different weights, which are displayed as follows:

```
==> netstat -rn

Destination      Gateway           Flags   Refs     Use  If   PMTU Exp Groups

Route tree for Protocol Family 2 (Internet):
default          9.3.149.65        UGc       0        0  en0     -   -   =>
default          10.10.10.3        UGc       0        0  en1     -   -
```

Here are two default routes through two different gateways: `9.3.149.65` and `10.10.10.3`. To configure weights and policies for these routes, as follows:

```
==> netstat -Cn
Routing tables
Destination      Gateway           Flags     Wt  Policy  If   Cost Config_Cost

Route tree for Protocol Family 2 (Internet):
default          9.3.149.65        UGc        2   LUT     en0     0      0 =>
default          10.10.10.3        UGc        4   -"-     en1     0      0
```

Route `1` through `9.3.149.65` is configured with a weight of `2` and a policy corresponding to `Lowest Utilization (4)`. Route `2` has a weight of `4`. The policy information is per multipath routing set, not individual routes.

To add route 1 and route 2, type the following commands:

```
route add default 9.3.149.65 -weight 2 -policy 4
route add default 10.10.10.3 -weight 4
```

Example 2: Changing Weight and Policy Information for Routes: To change weight and policy information for routes already created, use the following command:

```
route set
```

To change the weight and policy information of the default routes added in Example 1, type the following commands:

```
route set default 9.3.149.65 -weight 3 -policy 2
route set default 10.10.10.3 -weight 6
```

The output is as follows:

```
==> netstat -Cn
Routing tables
Destination      Gateway           Flags     Wt  Policy  If   Cost Config_Cost

Route tree for Protocol Family 2 (Internet):
default          9.3.149.65        UGc        3  RND     en0     0         0 =>
default          10.10.10.3        UGc        6   -"-    en1     0         0
```

Virtual SCSI client adapter

To gather problem determination information for virtual SCSI client adapters, run the `snap client_collect,all` command. The results are left in the **/tmp/ibmsupt/client_collect** directory.

The virtual SCSI client adapter can support up to 42 child devices simultaneously active without a performance degradation. Up to 84 devices can be configured on a single virtual SCSI client adapter, but performance will be less than optimal.

System management

AIX Network Data Administration Facility

The AIX Network Data Administration Facility (AIX NDAF) provides secure centralized management of file system data relationships across multiple systems. Its purpose is to facilitate central control of a federated file system namespace and replicated read-only data across a collection of AIX NFS Version 4 file servers. AIX NDAF resides on the AIX 5.3 Expansion Pack media.

Distributed Command Execution Manager (DCEM)

The installation of the CSM DCEM GUI (**csm.dcem.gui**) and the CSM DCEM Web-based System Manager application (**csm.dcem.websm**) packages is dependent upon the installation of the CSM Server. See “Cluster Systems Management” on page 21 for installation instructions.

Enhanced nimadm command

The **nimadm** command is enhanced to allow the system administrator to do the following:

- Use a NIM client’s **rootvg** to create a NIM **mksysb** resource that has been migrated to a new version or release level of AIX.
- Use a NIM **mksysb** resource to create a NIM **mksysb** resource that is migrated to a new version or release level of AIX.
- Use a NIM **mksysb** resource to restore to a free disk, or disks, on a NIM client and simultaneously migrate to a new version or release level of AIX.

Refer to the **nimadm** man page for further information and syntax.

Predefined XOPEN macros

The definition of the POSIX macros defined by the AIX system headers is as follows:

- **_ALL_SOURCE**, **_XOPEN_SOURCE** and **_XOPEN_SOURCE_EXTENDED**: checked by the system headers to enable/disable names provided by the runtime library
- **_ALL_SOURCE**: when defined, compiles for any strict standards mode

  It is set by default through `<standards.h>` if the strict standards scope feature macro is enabled. The **_ALL_SOURCE** macro enables all standards scopes and allows non-standard constructs to be visible as well.
- **_XOPEN_SOURCE**: when set to a value causes constructs defined in some UNIX specification to be visible

  When **_XOPEN_SOURCE** is defined, each header defines or declares some identifiers, potentially conflicting with identifiers used by the application. The set of identifiers visible to the application consists of precisely those identifiers from the header pages of the included headers, as well as additional identifiers reserved for the implementation. There are various values set to correspond to various levels of UNIX specifications:
  - **_XOPEN_SOURCE=1** is XPG4 or earlier
  - **_XOPEN_SOURCE=500** is UNIX98
  - **_XOPEN_SOURCE=600** is the correct feature macro to specify for strict UNIX03 support

  .
- **_XOPEN_SOURCE_EXTENDED**: is UNIX95

  Note: When **_XOPEN** was extended for UNIX98, this naming convention was dropped.

Appendix B. AIX 5L Version 5.3 unsupported devices

Unsupported devices and machines

The following devices and machines are not supported:

- RS/6000 or OEM hardware based on the MCA bus
- Scalable Parallel (SP™) nodes based on the MCA bus
- RS/6000, Power Personal Systems, or OEM hardware based on the PReP architecture
- POWER1, POWER2™, POWER Single Chip (RSC), POWER2 Single Chip (P2RSC), and 601 and 603 processors
- PCMCIA device support
- PCI adapters:
  - 2408 F/W SCSI SE, PCI/SHORT/32BIT/5V
  - 2409 F/W SCSI DIFF, EXT ONLY, PCI/SHORT/32BIT/5V
  - 2638 VIDEO CAPTURE (NTSC/PAL/SECAM), PCI/LONG/32BIT/5V
  - 2648 (GXT150P) PCI/SHORT/32BIT/5V, GRAPHICS ADAPTER
  - 2657 S15 GRAPHICS ADAPTER, PCI/SHORT/32BIT/5V, WEITEK P9100
  - 2708 Eicon ISDN DIVA PRO 2.0 PCI S/T Adapter
  - 2751 S/390 ESCON Channel PCI Adapter
  - 2837 MVP MULTI-MONITOR ADAPTER, PCI/LONG/32BIT/3.3 OR 5V
  - 2854 3D (GXT500P), PCI/LONG/32BIT/3.3 OR 5V, GRAPHICS ADAPTER
  - 2855 3DX (GXT550P), PCI/LONG/32BIT/3.3 OR 5V, GRAPHICS ADAPTER
  - 2856 PCI/SHORT/32BIT/3.3 OR 5V, 7250 ATTACH ADAPTER
  - 8242 10/100BASET ETHERNET PCI/SHORT/32BIT/5V
- ISA adapters:
  - 2647 VIDEO CAPTURE ENHANCEMENT, ISA/SHORT
  - 2701 4 PORT SDLC, ISA/LONG, EIA 232/V.35/X.21
  - 2931 8-PORT, ISA/LONG, EIA232 ADAPTER/FAN-OUT BOX
  - 2932 8-PORT, ISA/LONG, EIA232/422 ADAPTER/FAN-OUT BOX
  - 2933 128-PORT, ISA/LONG, EIA232 ASYNCH CONTROLLER
  - 2961 1 PORT X.25, SDLC, PPP, ISA/LONG, ADAPTER (C1X)
  - 2971 TOKEN RING ADAPTER, ISA
  - 2981 ETHERNET ADAPTER, ISA, RJ45/BNC
  - 8240 A/M 3COM ETHERNET ISA/SHORT TP ONLY
  - 8241 A/M 3COM ETHERNET ISA/SHORT BNC/AUI
- Non-CHRP Graphics Adapters:
  - Gt3/Gt3i
  - Gt4/Gt4e/Gt4i/Gt4x/Gt4xi
  - GXT110P
  - GXT150L/GXT150M/GXT150P
  - GXT155L
  - GXT500
  - GXT500D
  - GXT500P
  - GXT550P (FC 2855 only)

- GXT800M
- GXT1000™
- MVP MULTIPCI Adapter
- S15
- VIDEO OUTPUT OPTION (#3200) (FC 7254)
- 7250 ATTACH Adapter (FC 2856)

Unsupported functions and filesets

The following functions and filesets are not supported:

- 7318 Model P10/S20 Serial Communications Network Server
- AIX Xstation Manager®
- AIX Version 3.2 Network Installation Tools
- Remote Customer Support and Services
- SOMobjects® Base Toolkit
- Information Presentation Facility Runtime
- X11.vsm.helps
- X11.vsm.icons
- X11.vsm.rte
- GL 3.2
- power management
- IBM-850 locales
- libipfx.a
- devices.pci.b7105090
- The 7318 Serial Communications Network Server
- Network Terminal Accelerator
- The 9333 Serial Link DASD Subsystem
- devices.pci.331101e0
- OpenGL.html.xx_XX
- PEX_PHIGS.html.xx_XX
- X11.html.xx_XX
- bos.html.xx_XX.adapt
- bos.html.xx_XX.cmds
- bos.html.xx_XX.files
- bos.html.xx_XX.lowlevprg
- bos.html.xx_XX.manage_gds
- bos.html.xx_XX.prog_gds
- bos.html.xx_XX.techref
- bos.html.xx_XX.topnav
- bos.html.xx_XX.user_gds
- bos.man.xx_XX
- infocenter.html.xx_XX.tasks_topics
- perfagent.html.xx_XX
- sx25.html.xx_XX
- IMNSearch.bld
- IMNSearch.msg.xx_XX.rte.com

- IMNSearch.rte
- IMNSearch.rte.httpdlite
- devices.pci.14107800.rte
- devices.pci.esconCU.rte
- devices.common.IBM.esconCU.mpc.rte

Unsupported EEH devices

NOTE: THIS SECTION WILL BE REVISED WITH UNSUPPORTED INSTEAD OF SUPPORTED DEVICES

Device Driver support for Enhanced Error Handling (EEH) is limited to the following devices that are supported by AIX 5L Version 5.3:

- Storage Adapters:
  - PCI-X Dual Channel Ultra320 SCSI Adapter (5712, 5710, 1974)
  - PCI-X Dual Channel Ultra320 SCSI RAID Adapter (5703, 5711, 1975)
  - Dual Channel SCSI RAID Enablement Card (5709, 5726, 1976)
  - PCI-X Quad Channel U320 SCSI RAID Adapter (2780)
  - PCI-XDDR Dual Channel Ultra320 SCSI Adapter (5736, 1912)
  - PCI-XDDR Dual Channel U320 SCSI RAID Adapter (5737, 1913)
  - Dual Channel SCSI RAID Enablement Card (5727, 5728, 1907)
  - Dual Channel SCSI RAID Enablement Card (1908)
- Communications and connectivity (PCI bus type):
  - Token-Ring PCI 4/16 Adapter (FC 2920 and 4959)
  - IBM Ethernet 10/100 Mbps (FC 2968)
  - 10/100 Mbps Ethernet PCI Adapter II (FC 4962)
  - IBM 4-Port 10/100 Base-TX Ethernet PCI Adapter (FC 4961)
  - 10/100/1000 Base-T Ethernet PCI Adapter (FC 2975)
  - Gigabit Ethernet (FC 2969)
  - TURBOWAYS 622 Mbps PCI MMF ATM Adapter (FC 2946)
  - 2-Port Multiprotocol PCI Adapter (FC 2962)
  - 8-Port and 128-Port 232/422 Async PCI Adapters (FC 2943 and 2944)
  - IBM 64 bit/66 MHz PCI ATM 155 adapter (FC 4953 and 4957)
  - IBM Gigabit Ethernet-SX PCI-X Adapter (FC 5700)
  - IBM 10/100/1000 Base-TX Ethernet PCI-X Adapter (FC 5701)
  - IBM 2-Port 10/100/1000 Base-TX Ethernet PCI-X Adapter (FC 5706)
  - IBM 2-Port Gigabit Ethernet-SX PCI-X Adapter (FC 5707)
  - 10 Gigabit Ethernet-SR PCI-X Adapter (FC 5718)
  - 10 Gigabit Ethernet-LR PCI-X Adapter (FC 5719)
  - IBM ARTIC960HX 4-PORT PCI ADAPTER (FC 2947)
  - IBM ARTIC960RXD QUAD DIGITAL TRUNK ADAPTER (FC 6310)
  - 4-Port 10/100/1000 Base-TX PCI-X Adapter (FC 5740)
  - IBM 10 Gigabit Ethernet-SR PCI-X 2.0 DDR Adapter (FC 5721)
  - IBM 10 Gigabit Ethernet-LR PCI-X 2.0 DDR Adapter (FC 5722)
- Encryption Adapters:
  - IBM PCI Cryptographic Coprocessor (FC 4958 and 4963) *

  - IBM eBusiness Cryptographic Accelerator (FC 4960) *
- Graphics and Miscellaneous
  - GXT135P Graphics Adapter (FC 2848 and 2849) *
  - USB Open Host Controller (FC 2737 and 2738) *
  - GXT4500P (FC 2842)*
  - GXT6500P (FC 2843)*

Note: The devices above that are denoted with an asterisk (*) require the user to intervene and manually recover the device after a bus error is encountered (for example, through device reconfiguration). Also, you may need to reboot Graphics and USB devices because those devices may not completely recover. If the device encounters an error during the configuration process, the device will be left in the defined state until there is a subsequent configuration attempt.

Appendix C. Listing of filesets on the AIX media

AIX 5L for POWER Version 5.3 CD set

The AIX 5L for POWER Version 5.3 CD set consists of eight CDs with the following software groupings. During normal software installation with the CD device (such as **/dev/cd0**), you are prompted for the CD volume. Corresponding language filesets (messages and locales) are also installed by default, to match the environment of your system.

- **Volume 1:** Contains the minimal Base Operating System (BOS) software installed on every system, as well as all devices and both kernels, and the English message catalogs. To install a system with only ″Volume 1″ from the BOS menu, set Desktop to `NONE`, and change the default for Graphics bundle to `No`.
- **Volume 2:** Contains all the software to install the Graphics bundle, and to set the Desktop to `CDE`. System management software, X11 software and Java software are on this CD.
- **Volume 3:** Contains printer software, as well as additional CSM, RSCT, and Java software. The English man pages for libs and files, and additional software not installed by default, are on this CD.
- **Volume 4:** Contains pieces of Open_GL, PEX_PHIGS, LDAP, DB2 and the Fortran compiler software that comes with AIX.
- **Volume 5:** Contains software for installing AIX in other languages. Messages, help text, and locales for German, French, Italian and Japenese are included.
- **Volume 6:** Contains software for installing AIX in other languages. Messages, help text, and locales for Catalan, Czech, Spanish, Hungarian, Polish, Brazilian Portuguese, Russian, and Slovakian are included.
- **Volume 7:** Contains software for installing AIX in other languages. Messages, help text, and locales for Korean, Simplified Chinese, and Traditional Chinese are included.
- **Volume 8:** Contains software to support bidirectional and complex text languages, including Arabic, Hebrew, Thai, Vietnamese, Hindi, Tamil, Telugu, Gujarati, Marathi, Kannada, and Malayalam. Chinese language support for Hong Kong and Singapore is also included.

Appendix D. CAPP/EAL4+ updates

This appendix serves as a supplement to *AIX 5L Version 5.3 Security* and provides information relevant to configuring and using a system according to the requirements for the CAPP/EAL4+ Common Criteria evaluation of AIX 5L for POWER Version 5.3 with the 5300-05 Technology Level and the 5300-05-02 Service Pack (AIX 5300-05-02). To make this appendix consistent with *AIX 5L Version 5.3 Security*, the headings used are identical unless otherwise noted. The information here supersedes information in *AIX 5L Version 5.3 Security*.

For more information about CAPP/EAL4+, see *AIX 5L Version 5.3 Security* and AIX information.

CAPP/EAL4+ compliant system overview

AIX 5300-05-02 is evaluated on IBM pSeries Symmetric Multiprocessor systems for all System p5 servers with POWER5 and POWER 5+ processors.

Installing a CAPP/EAL4+ system

Any statements in *AIX 5L Version 5.3 Security* concerning AIX 5.2 are not applicable for this installation. After AIX 5300-05 is installed in CAPP/EAL4+ mode as described in *AIX 5L Version 5.3 Security*, you must upgrade to 5300-05-02.

The following steps describe how to download and install AIX 5300-05-02. Be sure to use P91209 and 5849 to download AIX 5300-05 and a copy of the release notes.

1. Use Download Director to securely download the filesets needed to upgrade to AIX 5300-05-02. Go to the Quick Links for AIX fixes web site at http://www-03.ibm.com/servers/eserver/support/unixservers/aixfixes.html.
2. In the **Search by** drop-down list, select **APAR number or abstract** and type `IY88827` in the text box.
3. Add the APAR to your download list by highlighting it and selecting **Add to my download list**.
4. Click **Continue**.
5. In the Packaging Options window, check the **Include prerequisites and corequisites** and **Include ifrequisites** packaging options. Do not check the **Include fixes that correct regressions** and **Replace superseded fixes with the latest** packaging options.
6. Select **5300-05** from the drop-down list.
7. Provide the output file for the **lslpp -Lc** command by selecting the **Browse** button and browsing to the location of the file.
8. Click **Continue**.
9. When the Download fixes window opens, select **Download all filesets using Java applet** to start the Download Director Java applet. You might need to grant the applet access to the system you are downloading the filesets to by responding to the pop-up dialog boxes in your browser.
10. Download and install the filesets using the Java applet. To install the filesets, place them in a directory on the system to be upgraded. In this example, the filesets are copied to the **/usr/sys/sp2** directory. Generate a **.toc** file by running the **inutoc** command:

    ```
    # inutoc /usr/sys/sp2
    ```

    After the **.toc** file is generated, run the following command to invoke smitty to install the updates:

    ```
    # smitty update_all
    ```
11. Run the **/usr/lib/security/CC_EVALify.sh** command.

Your system is now upgraded to AIX 5300-05-02. You must reboot the system and verify that the system has been upgraded by running the following command:

> # oslevel -r or oslevel -s

> If the system was successfully upgraded, **5300-05-02** is displayed.

>

> CAPP/EAL4+ system physical environment

> IPv6 is also included in the evaluated configuration, but only the functional capabilities of IPv6 that are
> also supported by IPv4 are included.

>

> CAPP/EAL4+ system configuration

> List of setuid/setgid programs

> • **/usr/bin/ipcs**
> • **/usr/bin/ipcs64**

> The **setuid** bit for the **ipcs** command should be removed by the system administrator. The system
> administrator should run the **chmod u-s /usr/bin/ipcs** and **chmod u-s /usr/bin/ipcs64** commands.

> Network configuration

> The statements are no longer applicable.

> System services

>

UID	Command	Description
root	**/usr/sbin/secldapclntd**	AIX LDAP authentication daemon
root	**/usr/sbin/gssd**	Services kernel requests for GSS operation
root	**/usr/sbin/nfsrgyd**	Name translation service for NFS v4 servers/clients

>

> Running a CAPP/EAL4+ distributed system

> The description of shared NFS is no longer applicable.

>

> NSF v4 Access Control Lists and contents policy

> An NFS v4 Access Control List (ACL) consists of a list of entries with the following fields:
> • The **Type** field contains one of the following values:
> – ALLOW – Grants the subject, specified in the **Who** field, the permission(s) specified in the **Mask** field.
> – DENY – Denies the subject, specified in the **Who** field, the permission(s) specified in the **Mask** field.
> • The **Mask** field contains one or more of the following fine grained permission values:
> – READ_DATA / LIST_DIRECTORY – Read the data from a non-directory object or list the objects in a
> directory.
> – WRITE_DATA / ADD_FILE – Write data into a non-directory object or add a non-directory object to a
> directory.
> – APPEND_DATA / ADD_SUBDIRECTORY – Append data into a non-directory object or add a subdirectory to
> a directory.
> – READ_NAMED_ATTRS – Read the named attributes of an object.
> – WRITE_NAMED_ATTRS – Write the named attributes of an object.
> – EXECUTE – Execute a file or traverse/search a directory.
> – DELETE_CHILD – Delete a file or directory within a directory.

> – READ_ATTRIBUTES – Read the basic (non-ACL) attributes of a file.
> – WRITE_ATTRIBUTES – Change the times associated with a file or directory.
> – DELETE – Delete a file or directory.
> – READ_ACL – Read the ACL.
> – WRITE_ACL – Write the ACL.
> – WRITE_OWNER – Change the owner and group.
> – SYNCHRONIZE – Synchronize access (exists for compatibility with other NFS v4 clients, but has no implemented function).

> • **Flags** field – This field defines the inheritance capabilities of directory ACLs and indicates whether the **Who** field contains a group or not. This field contains zero or more of the following flags:
> – FILE_INHERIT – Specifies that, in this directory, newly created non-directory objects inherit this entry.
> – DIRECTORY_INHERIT – Specifies that, in this directory, newly created subdirectories inherit this entry.
> – NO_PROPAGATE_INHERIT – Specifies that, in this directory, newly created subdirectories inherit this entry, but these subdirectories do not pass this entry to their newly created subdirectories.
> – INHERIT_ONLY – Specifies that this entry does not apply to this directory, only to the newly created objects that inherit this entry.
> – IDENTIFIER_GROUP – Specifies that the **Who** field represents a group; otherwise, the **Who** field represents a user or a special Who value.

> • **Who** field – This field contains one of the following values:
> – User – Specifies the user to whom this entry applies.
> – Group – Specifies the group to which this entry applies.
> – Special – This attribute can be one of the following values:
> - OWNER@ – Specifies that this entry applies to the owner of the object.
> - GROUP@ – Specifies that this entry applies to the owning group of the object.
> - EVERYONE@ – Specifies that this entry applies to all users of the system including the owner and group.

> If the ACL is empty, only a subject with an effective UID of 0 can access the object. The owner of an object implicitly has the following mask values regardless of what the ACL might or might not contain:

> • READ_ACL
> • WRITE_ACL
> • READ_ATTRIBUTES
> • WRITE_ATTRIBUTES

> The APPEND_DATA value is implemented as WRITE_DATA. Effectively, there's no functional distinction between the WRITE_DATA value and the APPEND_DATA value. Both values must be set or unset in unison.

> Object ownership can be modified through the use of the WRITE_OWNER value. When the owner or group is changed, the **setuid** bit is turned off. The inheritance flags only have meaning in a directory's ACL and only apply to objects that are created in the directory after the inheritance flags have been set (for example, existing objects are not affected by inheritance changes to the parent directory's ACL). The entries in an NFS v4 ACL are order dependent. To determine if the requested access is allowed, each entry is processed in order. Only entries that have the following values are considered:

> • A **Who** field that matches the effective UID
> • A user that is specified in the entry or effective GID
> • A group that is specified in the entry of the subject

> Each entry is processed until all of the bits of the requester's access have been ALLOWED. After an access type has been ALLOWED by an entry, it is no longer considered in the processing of later entries.

> If a DENY entry is encountered where the requester's access for that mask value is necessary and
> undetermined, the request is denied. If the evaluation reaches the end of the ACL, the request is denied.

> The maximum supported ACL size is 64 KB. Each entry in an ACL is of variable length and 64 KB is the
> only limit on an entry.

> The WRITE OWNER value

> The NFS v4 policy provides control over who can read and write the attributes of an object. A subject with
> effective UID 0 can always override the NFS v4 policy. The object owner can allow others to read and
> write the attributes of an object using the `READ_ATTRIBUTES`, `WRITE_ATTRIBUTES`, `READ_NAMED_ATTRS`, and
> `WRITE_NAME_ATTRS` attributes of the ACL mask. The owner can control who can read and write the ACL
> using the `READ_ACL` and `WRITE_ACL` values of the ACL mask. The object owner always has
> `READ_ATTRIBUTES`, `WRITE_ATTRIBUTES`, `READ_ACL`, and `WRITE_ACL` access. The object owner can also allow
> others to change the owner and group of the object using the `WRITE_OWNER` attribute. An object owner
> cannot change the owner or group of the object by default, but the object owner can add a `WRITE_OWNER`
> entry to the ACL specifying themselves, or the object can inherit an ACL entry that specifies a `WRITE_OWNER`
> entry with a Who value of `OWNER@`. When the owner or group is changed, the **setuid** bit is turned off.

> The following are some exceptions to the rules:

> • If the object is owned by UID 0, only UID 0 can change the owner, but the group can still be changed
> by a subject with the `WRITE_OWNER` attribute.

> • Assuming the object has the `WRITE_OWNER` attribute for the subject, in versions of AIX 5.3 prior to
> technology level 5300-05, if the object has a non-UID 0 owner, the owner can only be changed to
> another non-UID 0 user. In AIX technology level 5300-05 and later, if the object has a non-UID 0 owner,
> the owner can only be changed to the EUID of the subject attempting to change the owner.

> • The group can be changed to any group in the subject's concurrent group set with the exception that it
> can never be changed to GID 0 or GID 7 (system or security), even if these two groups are in the
> concurrent group set of the subject.

>

> LDAP-based and file-based administrative database supported

> This evaluation does not support NFS administrative database. Authentication methods such as DCE and
> NIS are not supported.

> It supports only the following:

> • File-based authentication (default)

> • UNIX-style LDAP-based authentication (use LDAP server ITDSv 6.0)

> The User Authentication section, under ″Securing the Base Operating System″ in *AIX 5L Version 5.3*
> *Security* gives the details on file-based authentication.

> For more information about file-based authentication, see User authentication in *AIX 5L Version 5.3*
> *Security*.

> LDAP authentication

> LDAP-base I&A is configured in the "UNIX-type" authentication mode. In this mode, the administrative data
> (including user names, IDs, and passwords) are stored in LDAP where access to the data is limited to the
> LDAP administrator. When a user logs into the system, the system binds to the LDAP server using the
> LDAP administrator account over an SSL connection, retrieves the necessary data for the user (including
> the password) from LDAP, and then performs authentication using the data retrieved from LDAP. The
> system maintains an administrative database on an LDAP server. The remaining hosts import the
> administrative data from the same LDAP server through the same mechanism previously described. The
> system maintains a consistent administrative database by making all administrative changes on the

> designated LDAP server. A user ID on any computer refers to the same individual on all other computers.
> In addition, the password configuration, name-to-UID mappings, and other data are identical on all hosts in
> the distributed system.

> For more information on LDAP authentication setup, see Light Directory Access Protocol. For more
> information in setting up SSL on LDAP, see Setting up SSL on the LDAP server and Setting up SSL on the
> LDAP client.

> LDAP server

> The **mksecldap -s** command sets up an AIX system as an LDAP server for security authentication and
> data management. Perform the following:

> • Use the RFC2307AIX schema with the **-S** option.

> • Set the server to use SSL by using the **-k** option. This requires installing the **GSKit** fileset and the
> **ldap.max_crypto_server** fileset. Use the **gsk7ikm** utility to generate the key pairs for the directory
> server.

> The LDAP user options must be set to satisfy the requirements of the evaluation. The RFC2370AIX
> schema defines the user attributes. Use the same values as described in CAPP/EAL4+ system
> configuration in *AIX 5L Version 5.3 Security*. The ITDS administrators are not forced to periodically change
> their passwords (for example, there's no MaxAge value for administrative passwords). Because of this, the
> LDAP administrative password must be changed as often as an AIX user (MaxAge = 8 (in weeks)).

> In ITDS 5.2, the authentication failure handling does not apply to Directory Administrator or to the
> members of the administrative group. Password composition rules also do not apply to administrative
> accounts. These need to be enforced if 5.2 is used.

> If the administrator does not use a common LDAP database backend for user management, the
> administrator must somehow ensure that the database that contains users credentials (listed below) is
> maintained consistently among the different TOE systems part of one network:

> **/etc/group**
> **/etc/passwd**
> **/etc/security/.ids**
> **/etc/security/.profile**
> **/etc/security/environ**
> **/etc/security/group**
> **/etc/security/limits**
> **/etc/security/passwd**
> **/etc/security/user**

> LDAP client

> The **mksecldap -c** command sets up an AIX system as an LDAP client for security authentication and
> data management. Perform the following:

> • Using the **mksecldap -c** command, specify **unix_auth** for the *authType* with the **-A** option.

> • Set the client to use SSL by using the **-k** option in the **mksecldap -c** command. Specifying the client
> SSL key requires installing the **GSKit** fileset and **ldap.max_crypto_client** fileset. Use the **gsk7ikm**
> utility to generate the key pairs for the directory server.

> For more information about LDAP, see the following:

> • Redbook: *Integrating AIX into Heterogenous LDAP Environments*.

> • Whitepaper: *Configuring an IBM Directory Server for User Authentication and Management in AIX*.

> • Whitepaper: *Configuring an AIX Client System for User Authentication and Management Through LDAP*.

NFS v4 Client/Server & Kerberos

The NFS v4 Client/Server environment includes LDAP for maintaining authentication data and Kerberos for establishing trusted channel between NFS v4 clients and servers. The evaluated configuration supports NAS v1.4 for Kerberos and ITDS v6.0 (LDAP server) for the user database. NAS v1.4 (Kerberos Version 5 Server) must be configured to use LDAP for its database. Kerberos tickets previously granted by the Kerberos server are valid until they expire.

When you are using Kerberos authentication, the credential used in remote procedure calls initiated by a user are associated with the current Kerberos ticket held by the user and is not influenced by the real or effective UID of the process. When you are accessing an NFS remote file system using Kerberos authentication while running a **setuid** program, the UID seen at the server is based on the Kerberos identity, not the UID that owns the **setuid** program being run.

The evaluated configuration involves setting up NFS to use RPCSEC-GSS security. For more information, see *Network File System*, *Configuring an NFS server*, and *Configuring an NFS client*. When setting up the server, choose Kerberos authentication and enable enhanced security on the server. You can enable this through SMIT using the **chnfs** command. The **chnfs** command has the option to enable RPCSEC_GSS security. When you are setting up the client, follow the instructions to use Kerberos in *Configuring an NFS client*. See *Setting up a network for RPCSEC-GSS* for the instructions to set up the Kerberos data server with DES3 encryption for security. The evaluated configuration supports only des3 encryption.

Note: This section does not currently exist in *AIX 5L Version 5.3 Security*

Password Rules

The evaluated configuration should have the following values for password rules when you are using the Kerberos server with LDAP as the database. For more information about password rules, see "Chapter 9. Managing Network Authentication Service passwords" in the *IBM Network Authentication Service Version 1.4 for AIX, Linux and Solaris Administrator's and User's Guide*.

≥ mindiff = 4

maxrepeats = 2

minalpha = 2

minother = 2

minlen = 8

minage = 0

histsize = 10

To have the AIX NFS v4 client and AIX NFS v4 server securely communicate explicitly using only DES3 enctypes, create the ″nfs/hostname″ server principal with DES3 enctype (such as **des3-cbc-sha1**), along with the corresponding entry in the **keytab** file (using **kadmin** interface) and have DES3 (such as **des3-cbc-sha1**) as the first entry in the **default_tgs_enctypes** section of the **/etc/krb5/krb5.conf** file on the NFS v4 client machine.

For more information about securing NFS, see *Securing NFS in AIX An Introduction to NFS v4 in AIX 5L Version 5.3*.

Virtual I/O Server

Virtual I/O Server (VIOS) resides in a separate LPAR partition and provides basic discretionary access control between VIOS SCSI device drivers acting on behalf of LPAR partitions and SCSI-based logical volumes and physical volumes through mappings. An LPAR partition (through a VIOS SCSI device driver) may be mapped to 0 or more logical and physical volumes, but a volume can only be mapped to one

> LPAR partition. This mapping limits an LPAR partition to only the volumes assigned to it. VIOS also
> controls the mapping of VIOS Ethernet adapter device drivers to VIOS Ethernet device drivers acting on
> behalf of groups of LPAR partitions sharing a virtual network. In the evaluated configuration, only a
> one-to-one mapping of an Ethernet adapter device driver to an Ethernet device driver acting on behalf of a
> group of LPAR partitions is allowed. The one-to-one mapping is configured by the administrator and
> enforced by the device drivers. Also, the Ethernet packets must not be tagged with a VLAN tag in the
> evaluated configuration. This mechanism can be used to limit which LPAR partitions see certain Ethernet
> packets.

> The VIOS interface should be protected from access by unprivileged users. The VIOS user options must
> be set to satisfy the requirements of the evaluation. The actual requirement is that the probability of
> correctly guessing a password should be at least 1 in 1,000,000 and the probability of correctly guessing a
> password with repeated attempts in one minute should be at least 1 in 100,000. The following parameters
> should be changed for the user in the **/etc/security/user** directory.

≥ maxage = 8
> maxexpired = 1
> minother = 2
> minlen = 8
> maxrepeats = 2
> loginretries = 3
> histexpire = 52
> histsize = 20
>

> To change the defaults, use the following commands:

```
> type oem_setup_env
>
> chsec -f /etc/security/user -s default  -a maxage=8 -a maxexpired=1 -a minother=2
>   -a minlen=8 -a maxrepeats=2 -a loginretries=3 -a histexpire=52 -a histsize=20
```

> When the prime administrator (**padmin**) creates a new user, the user attributes must be specified explicitly
> for that user. For example, to create a user with name *davis*, the **padmin** would use the following
> command:

```
> mkuser  maxage=8 maxexpired=1 minother=2 minlen=8 maxrepeats=2 loginretries=3
>   histexpire=52 histsize=20 davis
```

> The **padmin** should also stop the following daemons and then reboot.

> To remove **writesrv** and **ctrmc** from the **/etc/inittab** file:

```
>   sshd:      stopsrc -s sshd.
```

> To prevent the daemon from starting at boot time, remove the **/etc/rc.d/rc2.d/Ksshd** and
> **/etc/rc.d/rc2.d/Ssshd** files. After reboot stop the **RSCT** daemons:

```
> stopsrc -g rsct_rm
> stopsrc -g rsct
```

> All users, regardless of their roles, are to be considered as administrative users.

> The system administrator can run all of the commands except those in the following list that are limited to
> prime admin (**padmin**):

> • **chdate**
> • **chuser**

> • **cleargcl**
> • **de_access**
> • **diagmenu**
> • **invscout**
> • **loginmsg**
> • **lsfailedlogin**
> • **lsgcl**
> • **mirrorios**
> • **mkuser**
> • **motd**
> • **oem_platform_level**
> • **oem_setup_env**
> • **redefvg**
> • **rmuser**
> • **shutdown**
> • **unmirrorios**

> **Note:** This section does not currently exist in *AIX 5L Version 5.3 Security*

> X Server

> X Server should not be allowed to bind to port 6000. To prevent the X Server from binding (listening) on
> port 6000, edit the **xserverrc** file in the **/usr/lpp/X11/defaults** directory, and modify the `EXTENSIONS`
> variable to `EXTENSIONS="$EXTENSIONS -x abx -x dbe -x GLX -secIP"`.

Appendix E. Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Dept. LRAS/Bldg. 003
11400 Burnet Road
Austin, TX 78758-3498
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

AFS
AIX
AIX 5L
alphaWorks
AS/400
AT
AnyNet
C/370
CICS
CICS/MVS
CICS/VSE
CT
CUA
Common User Access
DB2
DFS
DirectTalk
developerWorks
eServer
ESCON
GDDM
HACMP
IBM
iSeries

xSeries
Lotus
Lotus Notes
MVS
Micro Channel
Micro-Partitioning
NetView
OS/2
OS/390
Operating System/2
PAL
POWER
POWER2
POWER3
POWER4
POWER5
POWER5+
POWERserver
Portmaster
PowerPC
PowerPC 601
PowerPC 603
PowerPC 604
PowerPC Architecture
PowerPC Reference Platform
Quietwriter
RACF
RETAIN
RS/6000
Redbooks
S/390
SAA
SP
System p
System p5
System/370
System/390
Tivoli
TotalStorage
VSE/ESA
VTAM
Wake on LAN
Xstation Manager

AltiVec is a trademark of Freescale Semiconductor, Inc.

Java and all Java-based trademarks and logos are registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Printed in U.S.A.

SC23-5201-02