AIX 5L Version 5.2

Release Notes

IBM

AIX 5L Version 5.2

Release Notes

Second Edition (May 2003)

Contents

Chapter 1. Read This Before Installation

Note: AIX 5L Version 5.2 Release Notes contains information that will aid you in the installation of AIX 5L Version 5.2. To view the most current version of the Release Notes, go to the online Release Notes in the IBM eServer pSeries Information Center. The Information Center is located at the following Web address:

```
http://publib16.boulder.ibm.com/pseries/en_US/infocenter/base
```

Required Hardware

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

As the root user, run the following to see if you have a CHRP machine:

```
bootinfo -p
```

To find out more about supported and unsupported items, see Appendix A, "AIX 5L Version 5.2 Supported and Unsupported Items", on page 39.

IBM eServer POWER4 pSeries Machines

Required Levels Hardware Management Console (HMC) and Firmware

The following items are required for using AIX 5L Version 5.2:

* Platform Firmware: Version 3 firmware or later.

To determine your platform firmware level, do the following:

1. At the command prompt, type:

   ```
   lscfg -vp | grep -p Platform
   ```

2. The last 6 digits of the ROM level represent the Platform Firmware date in the format, *YYMMDD*.

   Firmware with dates between 04/13/2002 and 10/24/2002 or has firmware level 2x*YYMMDD* is considered Version 2.

   Firmware with dates between 10/25/2002 and 03/24/2003 or has firmware level 3x*YYMMDD* is considered Version 3.

3. It is required that users that have installed AIX 5L Version 5.2 with the 5200-01 Recommended Maintenance package (APAR IY39795) also upgrade to the latest Version 3 firmware service pack, HMC: Release 3 Version 2.0.

The upgrade of Platform Firmware may cause some AIX Version 5.1 partitions to encounter problems when they reboot. To avoid these problems, apply APAR IY31961 on all AIX Version 5.1 partitions before upgrading the Platform Firmware.

For more information, see the following:

* HMC: See the *IBM Hardware Management Console for pSeries Installation and Operations Guide* in the Information Center.
* Platform Firmware: Contact your IBM Customer Engineer.
* AIX: See the *AIX 5L Version 5.2 Installation Guide and Reference* in the Information Center.

Known Limitations

* 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 maximum of 8 logical partitions which have a graphics adapter and USB adapter are allowed.

- The CPU Gard functions should be disabled if AIX 5L Version 5.2 and Platform Firmware levels older than October, 2002 are used together.

  For systems with Platform Firmware levels older than October, 2002 and AIX 5L Version 5.2, disable the CPU Gard function by typing the following:

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

  If Platform Firmware levels are upgraded, CPU Gard functions can be re-enabled. You can do this by typing the following:

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

  In either case, no system reboot is required for the changes to take effect.
- The eServer pSeries 690 model 681 (7040-681) supports a maximum system memory size of 512 GB with appropriate memory Feature Codes installed.

  AIX 5L Version 5.2 and Linux logical partitions can have nearly 512 GB logical partition memory sizes (503 GB after page table, hypervisor, and TCE table usage). AIX Version 5.2 and Linux logical partitions should have the **Small Real Mode Address Region** option chosen on the HMC partition profile memory panel, and must have it chosen for greater than 256 GB logical partitions.

Some AIX Systems May Not Boot From CD-ROM

Some AIX systems may not boot from CD-ROM because of firmware issues. To determine if your system will be affected by this problem, do the following (before migrating or installing a running 4.3 or 5.1 system):

1. At the command prompt, type:

   ```
   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.

Where to Find AIX 5L Version 5.2 Documentation

Information Center

The pSeries Information Center is an information portal for AIX and pSeries customers. From this site, you can access:

- AIX Version 4.3, 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/en_US/infocenter/base
```

AIX 5L Version 5.2 Installation Documentation on CD

AIX 5L Version 5.2 with the 5200-01 Recommended Maintenance package documentation is only available through the Information Center. It is not contained on the AIX documentation CD.

AIX 5L Version 5.2 documentation is contained on the AIX documentation CD. The documentation is made accessible by installing the documentation CD onto the hard disk, or by mounting the CD in the CD-ROM drive.

Note: For the most current documentation information, see the Web site.

Migration

Minimum Memory/Disk Requirements

The minimum memory required for AIX 5L Version 5.2 is 128 MB and a 2.2 GB hard drive.

pre_migration and post_migration Scripts

New **pre_migration** and **post_migration** scripts that are being shipped will run several commands to check if your system is ready for migration.

The **pre_migration** script saves all output in the **/home/pre_migration.***date* directory, and does the following:
* Verifies the system platform
* Lists the device filesets being removed
* Lists all other filesets being removed
* Lists the saved base configuration files that will not be merged
* Lists the configuration files that will be merged, and saves them in **/home/pre_migration.***date***/saved_configuration_files**
* Verifies fileset version consistency
* Creates a list of all filesets installed, to be used by the **post_migration** command
* Checks the size of the boot logical volume
* Checks that the major number of **rootvg** is 10
* Determines if Kerberos is being used
* Checks disk and memory sizes
* If migrating from Version 4.2, verifies the correct updates are applied

Please check the files in the **/home/pre_migration.***date* directory for conditions you may want to correct.

The **post_migration** script saves all output in the **/home/post_migration.***date* directory, and does the following:
* Verifies fileset version consistency.
* Checks the installation list from before the migration, and creates a list of filesets that did not move to a newer VRMF level (**/home/post_migration.***date***/filesets_to_be_updated**).
* Compares saved and merged configuration files, and saves the output in the **/home/post_migration.***date***/compared_saved_configuration_files** file.

The **pre_migration** script is on your product CD. To execute before migration, mount the CD and run the script.

To mount the CD, type the following:
```
# mount -v cdrfs -o ro /dev/cdn mountpoint
```

Where *n* is the number of the CD and *mountpoint* is the mountpoint.

The script will be *mountpoint***/usr/lpp/bos/pre_migration**, and can be copied to your system.

After you migrate, the **post_migration** script will be in the **/usr/lpp/bos** directory.

Migration from 4.2.1

A system running AIX Version 4.2.1 must be updated with the September 1999 or later update CD before migrating to AIX 5L Version 5.2. The CD label should have the number LCD4-0252-13 or later. To check a running system, verify that the **bos.rte.install** file is at level 4.2.1.17, or higher.

xlC.rte Fileset

If migrating to AIX 5L Version 5.2 from AIX Version 4.2.x or AIX Version 4.3.x, check the level of the **xlC.rte** fileset with 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 5L Version 5.2. Without APAR IY17981 installed, the migrated system could fail to boot.

APAR IY17981 is available from the following AIX Electronic Fix Distribution Site:

```
http://www.ibm.com/server/support
```

Trusted Computing Base (TCB)

Known Problems and Workarounds

/dev/dlcqllc: If you are on a migrated system with TCB enabled, you have the following error from `tcbck -n tree`:

```
3001-020 The file /dev/dlcqllc was not found
```

Run the following command to recreate the **/dev/dlcqllc** device:

```
# mkdev -c dlc -s dlc -t x25_qllc
```

This will recreate the **/dev/dlcqllc** device.

alt_disk_install Migration

The **nimadm** utility allows the user to use NIM and **alt_disk_install** to migrate an existing system from AIX Version 4.3 or AIX Version 5.1 to AIX Version 5.2. Refer to the **nimadm** command in *AIX 5L Version 5.2 Installation Guide and Reference* for more information.

bos.clvm.enh Fileset after Migration to AIX 5L Version 5.2

The **bos.clvm.enh** fileset is not installed when your system is migrating to AIX 5L Version 5.2. After migrating, users of the bos.clvm.enh fileset will need to re-install the fileset from AIX Version 5.2 installation media.

sysck Errors on X11.loc.nl_BE.Dt.rte

Users of the nl_BE language fileset might experience **sysck** error messages during migration to AIX 5L Version 5.2. If this is encountered, perform a force installation of the **X11.loc.nl_BE.Dt.rte** fileset.

xmodmap/Zh_CN Errors after Migration

After migrating to AIX 5L Version 5.2, an **lppchk -f** command might give the following warning message:

```
lppchk: 0504-206 File /usr/lpp/X11/defaults/xmodmap/Zh_CN/keyboard could not be located.
lppchk: 0504-206 File /usr/lpp/X11/defaults/xmodmap/Zh_CN could not be located.
```

If you have this condition, please reinstall the **X11.loc.Zh_CN.base.rte** fileset after migration.

KDE Desktop after Migration from AIX Version 4.3 to AIX 5L Version 5.2

If the CDE and KDE desktops are both installed on a migrated or installed AIX Version 5.2 system, the KDE desktop may not start up from the CDE login. To fix this problem, remove the following CDE startup from the **/etc/inittab** file:

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

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

Leave the following KDE entry in 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*.

When 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 installing the **bos.pmapi** fileset. If you do not do this, the machine will fail to boot when attempting to load the **pmtoolkit** fileset's kernel extension.

Check if 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.

The Performance Monitoring API is only supported on partitioned configurations using firmware level RH020305 and later.

Check your firmware level by typing the following at the command line:

```
lscfg -vp | grep -p Platform
```

SNMPv3

After migrating to AIX 5L Version 5.2, the SNMPv3 non-encrypted version will be running by default. If you had 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 that information, see ″Network Management″ in *AIX 5L Version 5.2 System Management Guide: Communications and Networks*.

Kerberos

In the past, the Kerberos library was statically linked to all the secure remote commands (**rlogin**, **telnet**, **rcp**, **rsh** and **ftp**). Starting in AIX 5L Version 5.2, all the secure remote commands will use the Kerberos Version 5 library and the GSSAPI library provided by the IBM Network Authentication Service Version 1.3. Network Authentication Service Version 1.3 is located on the expansion CD. The only LPP that is mandatory to install is the **krb5.client.rte** fileset. If you are migrating to AIX 5L Version 5.2 and had Kerberos Version 5 or Kerberos Version 4 installed, the install scripts will prompt you to install the **krb5.client.rte** fileset. Now the secure remote commands will support Kerberos clients and servers from both Native Kerberos 5 and DCE. For more information, see ″Understanding the Secure Rcmds″ in the *AIX 5L Version 5.2 System User's Guide: Communications and Networks*.

AIX Toolbox for Linux Application Migration Information

Note: AIX 4.3.3 Toolbox applications will not run correctly on AIX 5L Version 5.2 unless the steps in this section are performed.

If you previously installed the AIX Toolbox for Linux Applications and the level of the **rpm.rte** fileset is less than 3.0.5.20, remove that software from the system before migrating to AIX 5L Version 5.2. 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 5.2 because of shared library restructuring.

Also, 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 RPMS, use the **destroyrpms** tool available in the **/tools** directory on the AIX Toolbox for Linux Applications CD by typing the following:

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

Additional information is also available on this CD in the **/README.TXT** file.

Multi-path I/O (MPIO)

After upgrading to AIX 5L Version 5.2 with the 5200-01 Recommended Maintenance package, some disk devices will no longer be configured as *other FC disk*. These devices instead will be configured as *MPIO other FC disk*. The devices that will be affected are EMC SYMMETRIX, HDS OPEN, and SSG SHARK disk subsystems. These devices will be configured as an MPIO device if the device was previously configured as *other FC disk*.

The following describes some of the similarities and differences that will be 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 predefines are installed.
- Is meant to be a transitory state during boot and install. The vendor-specific device ODM predefines 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 predefines are installed.
- Presently are not supported by PowerPath, MDS or SSD path management products. To support any of these products the vendor-specific no-MPIO ODM predefines 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 predefines 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 deinstalled with the force option, the AIX Version 5.2 release of MPIO will handle the migration. If the system is rebooted, the device instance will be in the define state. During migration, the device instance will be left in the define state and a new *other FC disk* instance will be created.

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

If the system is not rebooted and the device instance is in the available state, the device instance will be left unchanged. There may also be *other FC device* instances created. If the *MPIO other FC device* is not open, there will be an *other FC device* instance created for each path the device is detected on. If the *MPIO other FC device* is in the open state, there will not be any *other FC device* instances created. This is because the *MPIO other FC device* will have already issued a **SCIOSTART** command to the FC adapter for each of the paths. The FC adapter will 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* may cause indeterminate device behavior or data corruption. It is recommended that the system be rebooted to clear this condition. Once the system is rebooted, the *MPIO other FC device* instance will be in the defined state and can be removed using the **odmdelete** command. Running the **rmdev** command will not remove the device due to the missing predefines.

See "Multi-path I/O (MPIO)" on page 32 in these release notes for additional features of MPIO devices.

Installation

This section contains information about installing AIX 5L Version 5.2 that supplements the information contained in the AIX 5L Version 5.2 installation documentation.

The following publications describe AIX 5L Version 5.2 installation:
- *AIX 5L Version 5.2 Operating System Installation: Getting Started*
- *AIX 5L Version 5.2 Installation Guide and Reference*

Both installation guides are available online in the pSeries Information Center under AIX 5L Version 5.2 documentation 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. Specify the order number of the book you want to order.

To obtain AIX 5L Version 5.2 installation hints and tips, send e-mail with a subject line of **52_Install_Tips** to aixserv@austin.ibm.com, the AIX Service Mail Server.

Installing AIX 5L Version 5.2

The following methods can be used to install AIX 5L Version 5.2:

- Complete overwrite installation
- Preservation installation
- Migration installation

Note: After installing or migrating a system to AIX 5L Version 5.2, 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 5.2 to a lower level of AIX are not supported.

To install AIX 5L Version 5.2, boot from the product media and follow the instructions in the *AIX 5L Version 5.2 Installation Guide and Reference*, order number SC23-4389.

Note: AIX 5L Version 5.2 cannot be installed on MCA (Micro Channel Architecture) or PowerPC Reference Platrorm (PReP) machines.

For more information on migration, see "Migration" on page 3.

Software License Agreements (SLA)

When using ASCII displays, there are some instances where the Software License Agreements may not be fully displayed. In this event, the License Agreements can be viewed in all languages from the following link:

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

Portions of the pt_BR Portuguese locale were improperly translated, preventing the installation of AIX to complete. The pt_BR user should first install AIX with the C locale. Then, install the pt_BR locale and set the locale to pt_BR using SMIT.

New Base Operating System Installation Menus

When performing a base operating system (BOS) installation, if you select **Change/Show Installation Settings and Install** from the Welcome to Base Operating System Installation and Maintenance menu, the third menu option has changed from **Advanced Options** to **More Options**. If you select **More Options**, the **Install Options** menu displays the following choices on 32-bit systems:

```
1.Desktop.........................................CDE
2.Enable Trusted Computing Base...................No
3.Graphics Software..............................Yes
4.Documentation Services Software................No
5.Enable System Backups to install any system.....Yes
```

Note: The menu choices change depending on whether the console is graphical or not, and whether the system is 64-bit capable or not. The options also vary based on the type of installation.

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 will be slightly different:

SCSI disks	Disks may be grouped by adapter or scsi bus.
IBM TotalStorage FAStT Storage Servers	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.

Graphics Systems
The Desktop choices are CDE, KDE, GNOME, and NONE. If you select NONE, the following are still installed with the Graphics Bundle:

- X11
- Java
- Web-based System Manager

The Graphics bundle installation always defaults to yes (to be installed), whether the console is graphical or not. It allows for remote access to the applications from other graphical systems. The Desktop choice is not available if the console is non-graphical.

Enable 64-bit Kernel and JFS2
Beginning in AIX 5L Version 5.2, the following are now two separate choices:

- Enable 64-bit Kernel
- Create JFS2 File Systems

If the system is 64-bit, you may select to just have the 64-bit kernel as the running kernel, and create JFS file systems, or select to have both the 64-bit kernel and JFS2 file systems. If JFS2 file systems are created, the kernel must be 64-bit.

These options are only available for new and complete overwrite installations. During migration or preservation installations, the kernel and file systems will be retained as is.

Memory Requirements

AIX 5L Version 5.2 requires a minimum of 128 MB of physical memory.

Paging Space Requirements

AIX 5L Version 5.2 creates a 512 MB paging space (**/dev/hd6**) for all new and complete overwrite installations.

Disk Space Requirements

AIX 5L Version 5.2 requires a minimum of 2.2 GB of physical disk space for the same set of installed filesets due to the increased library sizes and additional function.

Note: The following measurements provide information on disk usage when you install AIX 5L Version 5.2 as compared to previous versions.

Base AIX Installation (Graphical System with CDE-Default)

	AIX 4.3.3 Allocated (Used)	AIX 5L for POWER Version 5.1 Allocated (Used)	AIX 5L Version 5.2 with the 5200-01 Recommended Maintenance package Allocated (Used)
/	4MB (2.5MB)	8MB (5.6MB)	16MB (9MB)
/usr	294MB (279MB)	385MB (370MB)	696MB (672MB)
/var	4MB (1.3MB)	4MB (1.4MB)	16MB (5MB)
/tmp	16MB (0.6MB)	20MB (0.9MB) see note	24MB (0.8MB) see note
/opt	N/A	4MB (0.2MB)	24MB (7MB)

Note: If **/tmp** is less than 32 MB, it will be increased to 32 MB during a migration installation so that the AIX 5L Version 5.2 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 logical volume and file system (**/dev/hd10opt** and **/opt** respectively), will be created and the data that existed in the **/opt** directory will be moved to the new file system.

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

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

Fixes for Installing the RH021019 Level of the pSeries 670 or pSeries 690 Firmware

This section describes fixes that should be considered when installing the RH021019 level of the pSeries 670 or pSeries 690 (7040-671 or 7040-681) firmware.

The following fixes are recommended prior to installing the RH021019 level of the p670 or p690 firmware:
- HMC version update
- AIX 5L Version 5.2. APAR
- AIX 5L for POWER Version 5.1 APARs

Before installing the RH021019 level of the p670 or p690 (7040-671 or 7040-681) firmware, upgrade the HMC to HMC Update Release 3 Version 1.0. After upgrading to HMC Update Release 3 Version 1.0, apply the HMC fix level *HMC Update Release 3 Version 1.2*. HMC support and HMC Update Release 3 Version 1.2 is available from the following Web sites:

HMC support page:

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

HMC PTF page:

`http://techsupport.services.ibm.com/server/hmc/corrsrv.html`

AIX support and APARs can be obtained at the following Web site:

`http://techsupport.services.ibm.com/server/support?view=pSeries`

If AIX 5L Version 5.2 or AIX 5L for POWER Version 5.1 with the 5100-03 Recommended Maintenance package is being used, an AIX fix to **csm.client** is recommended for DLPAR operations and the Service Focal Point operation. The APAR numbers for the **csm.client** fix are as follows:
- AIX 5L Version 5.2: IY34493
- AIX 5L for POWER Version 5.1 with the 5100-03 Recommended Maintenance package: IY34707

The **csm.client** fix addresses a problem where a configuration daemon relating to the HMC function used in DLPAR operations and Service Focal Point operation eventually stops, preventing future DLPAR operations and Service Focal Point operations.

If AIX 5L for POWER Version 5.1 with the 5100-02 Recommended Maintenance package is being used, an AIX fix to the AIX boot image is strongly recommended. It is recommended that this AIX fix be applied to all instances of the AIX 5L for POWER Version 5.1 with the 5100-02 Recommended Maintenance package before installing the RH021019 level of the p670 or p690 (7040-671 or 7040-681) firmware. The APAR number for this fix (IY31421) is available at the AIX support Web site given above.

APAR IY31421 corrects a boot failure (system may not boot, may hang or crash during the AIX boot process) that can occur with the RH021019 level of the p670 or p690 (7040-671 or 7040-681) firmware.

If the boot failure occurs because APAR IY31421 was not applied, do the following:

1. Boot the AIX installation CD-ROM.
2. Apply the IY31421 APAR.
3. Reboot the AIX disk image.

APAR IY31421 is only necessary for AIX 5L for POWER Version 5.1 with the 5100-02 Recommended Maintenance package. All other (earlier and later) recommended maintenance packages of AIX do not require APAR IY31421.

To determine the AIX recommended maintenance package level, type `oslevel -r AIX`. The command will return 5100-02, 5100-03, and so on.

Summary
The following is a summary of the recommended procedures:

- If you have AIX 5L Version 5.2, update images with APAR IY34493. This can also be done after the HMC and firmware updates.
- If you have AIX 5L for POWER Version 5.1 with the 5100-03 Recommended Maintenance package, update images with APAR IY34707. This can also be done after the HMC and firmware updates.
- If you have AIX 5L for POWER Version 5.1 with the 5100-02 Recommended Maintenance package, update images with APAR IY31421. This is strongly recommended as the first step.

After you have updated the appropriate level of AIX, perform the following tasks:

1. Update the HMC to the latest level listed on the HMC Web site.(`http://techsupport.services.ibm.com/server/hmc`)
2. Update the p670 or p690 firmware to the RH021019 level.

Other Installation Information

Installation Packaging Formats
AIX 5L Version 5.2 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

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

The AIX 5L Version 5.2 product media contains **installp** packages and RPM packages that are installed during a BOS installation. The **installp** packages are located in the following path:

/mount_point/installp/ppc

The RPM packages are located in the following path:

/mount_point/RPMS/ppc

If you have media that contains ISMP packages for AIX 5.2, 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 on software packaging, see the Software Product Packaging Concepts section in the *AIX 5L Version 5.2 Installation Guide and Reference*.

Example:
If you are using the **bffcreate** command or a user interface to copy images from CD to a directory on the disk, new subdirectories will be created in the target directory and the images will be placed in those directories.

In AIX 4.3, the default target directory for the **bffcreate** command was the **/usr/sys/inst.images** directory.

In AIX 5L Version 5.2, if the following command is used on a PowerPC machine:

```
bffcreate -d /dev/cd0 all
```

the default target directory is the **/usr/sys/inst.images/installp/ppc** directory.

Use the **gencopy** command as follows:

```
gencopy -d /dev/cd0 all
```

In addition to all the **installp** images, you would also get the **rpm** images that are contained on the product media, as follows:

```
cd /usr/sys/inst.images/RPMS/ppc
find . -print
./cdrecord.aix4.3.ppc.rpm
./mkisofs.aix4.3.ppc.rpm
```

Because of this change, you may have to update any custom scripts in which image locations are hardcoded.

DVD-RAM Drive Support for System Backup
In AIX 5L Version 5.2, support was added to create universal disk format (UDF) backups to DVD-RAM media. System backups (made with the **mkcd** command) have a limitation in that they expect the media to be 4.7 GB or larger per side. The **mkcd** command will not process the next volume until it writes over 4 GB on the current volume, thus the use of smaller media would result in corruption when going beyond the media's capacity. The other commands (such as **tar**, **cpio**, **backup**) do not have this size constraint.

For more information on creating CD and DVD backups, refer to the **/usr/lpp/bos.sysmgt/mkcd.README.txt** file.

To support booting from UDF backup media, the system requires the latest firmware. For firmware updates, refer to the following Web site:

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

The alt_disk_install Command and Dynamic Logical Partitioning
It is possible, on a system that supports dynamic logical partitioning (DLPAR), to dynamically add an adapter with disks to a running logical partition (LPAR). A new **rootvg** can then be installed to these added disks using the **alt_disk_install** command (**clone** or **mksysb**). There are a few flags to keep in mind if you are running the **alt_disk_install** command with dynamically added target disks on an LPAR system:

- If the target disk will be used to boot an LPAR other than the one where the operation is being executed, the **-O** (device reset) flag should be used.
- To prevent **bootlist** from being executed, the **-B** flag should be added. A general limitation of dynamically-added disks is that you can not specify them as a boot device (before an initial reboot). If you are attempting to boot an LPAR from dynamically added disks, the bootlist should be set in SMS menus.
- Dynamically-added disks will not appear bootable to the operating system (before an initial reboot). The user will need to verify that the newly added adapter and disks are bootable and run the **alt_disk_install** command with the **-g** (ignore bootability checks) flag.

BOS Installation Support for Fibre Channel Boot

BOS installation supports installation to Fibre Channel-attached disks. Either bootable AIX 5L Version 5.2 installation media or NIM resources prepared from such media is required. The Fibre Channel-attached disks must be attached to a Fibre Channel host adapter that supports boot capability. However, Fibre Channel drivers that have multiple, physical connections to the host or that require supplemental device software should not be chosen as **rootvg** drives.

A Fibre Channel-attached disk can be identified by a World Wide Port Name and Logical Unit ID. To see the format of the World Wide Port Name and Logical Unit ID, type:

```
lsattr -E -O -l DiskName
```

- In a non prompted BOS installation, you can specify a Fibre Channel-attached disk in the target_disk_data stanza of the **bosinst.data** file using the following as an example:

  ```
  SAN_DISKID = (World Wide Port Name)//(Logical Unit ID)
  ```

  In the above example, *(World Wide Port Name)* and *(Logical Unit ID)* are each in the format returned by the **lsattr** command, that is, ″0x″ followed by 1-16 hexadecimal digits.

- In a prompted BOS installation, the BOS menus display the list of available disks and associated information. You can select the desired disks.

System Dump

The default system dump device is paging space. This default may not be adequate for system configurations that have large memory sizes or when system availability is a concern.

For large memory systems, AIX allocates a dedicated dump device (**/dev/lg_dumplv**) if disk space is available. The following table shows the sizes of the dump devices based on the amount of real memory present in the system.

Dump Device Size (Based on System Real Memory)

Real Memory	Dump Device
4 GB < 12 GB	1 GB
12 GB < 24 GB	2 GB
24 GB < 48 GB	3 GB
>= 48 GB	4 GB

Note: If the root user starts a dump to a dump device set to **/dev/sysdumpnull**, the system will crash and no dump will occur.

The dump is now copied to tape in **pax** format instead of **tar**, because **pax** supports large (greater than 2 GB) files.

The **dumpcheck** facility runs by default each day at 3:00 p.m. local time. See the root **crontab** file to change the time that **dumpcheck** runs. The **dumpcheck** facility ensures that there is enough space in the dump device and copy directory to receive the system dump. If dump compression will facilitate a system dump, it will be enabled unless you have manually turned it off. The results of **dumpcheck** are in the system error log.

Network Installation Management

The 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 5.2 NIM product and includes the following topics:

- Restrictions on SPOT Creation for Releases Prior to 5.2 (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

Reliable Scalable Cluster Technology (RSCT)

The 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 at **/usr/sbin/rsct/README/rsct.core.README**, and contains additional information about the RMC application.

Restriction for Japanese Locales: When the responses specified with the predefined **notifyevent** script are used in a 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 corrupted. To work around this problem, you can modify the **notifyevent** script to not use the **$ERRM_COND_NAME** environment variable in the mail subject.

Firmware Limitation

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`.

The 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.

The problem is that the firmware on the affected machines cannot read the boot image from the hard disk if any part of the boot image resides past 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 the better 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*

   Replace *bootlv00* with 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.

As a reminder, this problem does not occur unless you have created and initialized a boot image that resides above the 4 GB boundary.

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
```

Emergency Fix Management

Emergency Fix (efix) Management provides utilities that package, install, and manage efixes. This involves the following commands:

epkg   The efix packager.

emgr   The efix manager.

The **epkg** command creates efix packages that can be installed by the **emgr** command. Once the efix is installed, the **emgr** command can be used to list, check, remove, and perform other operations with the installed efixes.

Service

Fixes and Problem-Solving Databases

You can download AIX fixes from the following IBM Server Support Web site:

```
http://www.ibm.com/server/support
```

You can also search technical databases, including:

- APARS
- Tips for AIX administrators

Chapter 2. AIX 5L Version 5.2

The following are some of the features that were introduced in AIX 5L Version 5.2:

- Dynamic Logical Partitioning (DLPAR)
- Capacity Upgrade on Demand (CUoD)
- Advanced RAS features
- Enhancements to Workload Manager
- Cluster Systems Management (CSM) for monitoring and administering multiple machines (both AIX and Linux) from a single point of control

To find out more about the new features introduced in AIX 5L Version 5.2, refer to the Information Center at the following Web site:

http://publib16.boulder.ibm.com/pseries/en_US/infocenter/base

Base Operating System (BOS)

Dynamic Logical Partitioning (DLPAR)

DLPAR provides great flexibility in the management of logical partitions. You can add, remove, or move processors, memory, and PCI slots between partitions as capacity requirements change without having to reboot any of the affected partitions, resulting in higher levels of resource utilization. Also, with this new capability you can establish separate test and production partitions and, over time, transition the resources from the production to the test partition, in effect creating a new production environment without experiencing any down time.

It is easy to setup your system for DLPAR by doing the following:

1. Install AIX 5.2.
2. Configure the network.
3. Reboot AIX.
4. Wait five minutes.

You are now ready to perform DLPAR operations at the HMC.

One requirement is that each logical partition must have its own network adapter

Logical partitions may be installed with either AIX 5.1 or 5.2, although to perform DLPAR operations, it is necessary to use AIX 5.2 in all of the affected partitions.

The dynamic reconfiguration of processors and memory is well integrated into AIX. In general, there are no application reconfiguration issues that need to be addressed. On the other hand, the dynamic reconfiguration of PCI slots does require the customer to log into the affected partition(s) and perform the relevant Hot Plug Procedures before you move or remove the DLPAR and after DLPAR add or move commands.

Adapter Hot Plug is an existing AIX feature. Use the search key words ″hot plug″ in AIX online documentation.

Both the 32-bit and 64-bit kernels support dynamic partitioning. By design, the 32-bit kernel supports only 96 GB of memory, so if you want to use more memory, then you should use the 64-bit kernel.

Set your minimum and maximum CPU and memory requirements in the LPAR's profile. To change the minimum memory less than desired, check the Small Real Mode Address Region box under Memory tab in the LPAR profile dialog.

For detailed instructions on executing DLPAR adapter actions, refer to the *eServer pSeries Planning for Partitioned-System Operations* book in the Information Center.

Asynchronous I/O (AIO)

Legacy AIO has a new subroutine called **aio_nwait_timeout** (**_posix_aio_nwait_timeout** in POSIX). It extends the capabilities of the **aio_nwait** subroutine by adding a fourth parameter to specify timeout values. It is not available through the **/usr/include/aio.h** file or **libc.a** library, so it must be accessed directly from the kernel.

The **aio_nwait_timeout** subroutine is a new system call that is exported from the AIO subsystem and is declared as the following:

```
int aio_nwait_timeout (max_cnt, req_cnt, aiocbpa, timeout)
int max_cnt;
int req_cnt;
struct aiocbp **aiocbpa;
uint timeout;
```

The first three arguments of this subroutine are exactly the same as the first three arguments of the **aio_nwait** subroutine. The **aio_nwait_timeout** subroutine waits for a certain number of asynchronous I/O operations to complete as specified by the *req_cnt* parameter, or until the call has blocked for a certain duration specified by the *timeout* parameter.

The *max_cnt* parameter, indicates the maximum number of pointers to the **aiocbp** structure that can be copied into the **aiocbpa** array.

The timeout value is specified in units of milliseconds. A timeout value of -1, indicates that the subroutine should behave like the **aio_nwait** subroutine, blocking until all of the requested I/O operations complete or until there are no more asynchronous I/O requests pending from the process.

A timeout value of 0, indicates that the subroutine will return immediately with the current completed number of asynchronous I/O requests. All other positive timeout values indicate that the subroutine needs to block until either the timeout value is reached or the requested number of asynchronous I/O operations complete.

The **aio_nwait_timeout** subroutine has to be redefined and accessed from POSIX AIO as **_posix_aio_nwait_timeout**.

To access the subroutine in Legacy AIO when compiling, use the **aio_nwait_timeout.imp** import file, which contains the following:

```
#!/unix

aio_nwait_timeout
```

Assuming the **aio_nwait_timout** subroutine is called inside of the **testcase.c** file, compile it with the following:

```
cc -D_AIO_AIX_SOURCE -o testcase -bI:./aio_nwait_timeout.imp  testcase.c
```

To access the subroutine in POSIX AIO when compiling, use the **aio_nwait_timeout.imp** import file which contains the following:

```
#!/unix

_posix_aio_nwait_timeout
```

Assuming the **_posix_aio_nwait_timeout** subroutine is called inside of the **testcase.c** file, compile it with the following:

```
cc -o testcase -bI:./aio_nwait_timeout.imp testcase.c
```

64-bit Kernel

AIX 5L Version 5.2 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.

Note: The 32-bit kernel continues to be supported on AIX 5L Version 5.2. The maximum real memory supported by a 32-bit kernel system (or partition) is 96 GB.

System Support

For information about supported and unsupported items, see Appendix A, "AIX 5L Version 5.2 Supported and Unsupported Items", on page 39.

Base Functionality

For the most part, the AIX 5L Version 5.2 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 functions:

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 5.2 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 may have to be recompiled.

Application Scalability: AIX 5L Version 5.2 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 5.2. In addition, existing 32-bit kernel extensions and device drivers used by 64-bit applications may 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 5.2 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 5.2 base media and are installed on all hardware systems that they support. By default, the 32-bit kernel is enabled at base system installation. However, you can override this at installation time to enable the 64-bit kernel through the system installation panels. (See "New Base Operating System Installation Menus" on page 8 for details on enabling the 64-bit kernel at BOS installation time.)

It is possible to switch between the 32-bit and 64-bit kernels without reinstalling the operating system. This can be done by modifying the **/usr/lib/boot/unix** and **/unix** directories to be a symbolic link to the binary for the desired kernel, running the **bosboot** command to write a new system boot image, and rebooting. The path name of the 64-bit kernel is **/usr/lib/boot/unix_64**, while the path names of the uniprocessor and multiprocessor versions of the 32-bit kernel are **/usr/lib/boot/unix_up** and **/usr/lib/boot/unix_mp**.

JFS2

JFS2 (enhanced journaled file system) is a file system, introduced in AIX 5L for POWER Version 5.1, that provides the capability to store much larger files than the existing Journaled File System (JFS). Customers can choose to implement either JFS, which is the recommended file system for 32-bit environments, or JFS2, which offers 64-bit functionality.

Note: Unlike the JFS file system, the JFS2 file system will not allow the **link** subroutine to be used on files of **directory** type. This limitation may cause some applications that operate correctly on a JFS file system to fail on a JFS2 file system.

The following table provides a summary of JFS and JFS2 functions:

Functions	JFS2	JFS
Fragments/Block Size	512-4096 Block sizes	512-4096 Frags
Maximum file system size	16TB	1TB
Maximum file size	16TB	64GB
Number of i-nodes	Dynamic, limited by disk space	Fixed, set at file-system creation
Directory organization	B-tree	Linear
Compression	No	Yes
Quotas	No	Yes

Note: The maximum file size and maximum file system size is limited to 1 Terabyte when used with the 32-bit kernel.

Data Management Application Programming Interface (DMAPI)

AIX 5L Version 5.2 with the 5200-01 Recommended Maintenance package provides a new Data Management Application Programming Interface (DMAPI), an implementation of the X/Open standard, *System Management: Data Storage Management (XDSM) API* published by The Open Group. The DMAPI allows software vendors to develop data management applications using a set of functions and semantics not found in POSIX-compliant systems. It does not provide direct functionality to the end user. Complete documentation of the DMAPI can be found in the Publications section of The Open Group's Web site (http://opengroup.org).

The DMAPI is implemented in an abstract layer within AIX, allowing any underlying file system to define its individual level of support and implementation options. The JFS file system does not provide any support for the DMAPI. The JFS2 behaviors for implementation options, limits, and other specifics described by the X/Open standard are outlined in the following paragraphs.

The **dm_init_service** subroutine returns 0 when the AIX DMAPI is properly initialized. It returns -1 if the initialization fails. Use of any other DMAPI subroutine after initialization fails will also fail.

The **dm_get_config** subroutine returns the following values for the JFS2 implementation options and limits:

DM_CONFIG_BULKALL	supported
DM_CONFIG_LEGACY	supported
DM_CONFIG_PERS_ATTRIBUTES	supported
DM_CONFIG_PERS_EVENTS	supported
DM_CONFIG_PERS_INHERIT_ATTRIBS	supported
DM_CONFIG_PERS_MANAGED_REGIONS	supported
DM_CONFIG_PUNCH_HOLE	supported
DM_CONFIG_WILL_RETRY	supported
DM_CONFIG_CREATE_BY_HANDLE	not supported
DM_CONFIG_LOCK_UPGRADE	not supported
DM_CONFIG_OBJ_REF	not supported
DM_CONFIG_PENDING	not supported
DM_CONFIG_DTIME_OVERLOAD	TRUE
DM_CONFIG_MAX_ATTR_ON_DESTROY	128
DM_CONFIG_MAX_ATTRIBUTE_SIZE	4072
DM_CONFIG_MAX_HANDLE_SIZE	32
DM_CONFIG_MAX_MANAGED_REGIONS	167
DM_CONFIG_MAX_MESSAGE_DATA	65536
DM_CONFIG_TOTAL_ATTRIBUTE_SPACE	4072

For AIX, **DM_SESSION_INFO_LEN** is 256, and **DM_ATTR_NAME_SIZE** is 8.

In this release, all DM attribute values share the same allocation. Consequently, the size of any one attribute's value cannot exceed **DM_CONFIG_MAX_ATTRIBUTE_SIZE** and is further restricted by the sum of the value sizes of all DM attributes associated with an object, which is also limited to **DM_CONFIG_MAX_ATTRIBUTE_SIZE**.

AIX DMAPI does not provide support for the **dm_downgrade_right**, **dm_upgrade_right**, the **dm_obj_ref_*** family, or **dm_pending** optional DMAPI subroutines. The JFS2 implementation does not support the optional DMAPI events cancel and debut, nor the additional optional **dm_getall_dmattr**, **dm_create_by_handle**, and **dm_symlink_by_handle** subroutines.

When a data management (DM) application specifies that it wants to block until a right becomes available, the DM application is blocked and can not be interrupted.

AIX allows multiple, non-overlapping persistent managed regions. For JFS2, only regular files are allowed to have managed regions. Due to the current implementation of JFS2's extended attribute support, the **dm_set_region** subroutine causes the file's *ctime* to be modified. JFS2 does not attempt to reorder nor coalesce managed regions.

When no session has registered to receive a particular event for which an object is enabled, and activity occurs which would otherwise trigger the event, AIX does not generate the event, and allows the process to proceed as if there is no event enabled.

Calling the **dm_set_eventlist** subroutine causes a persistent event list to be stored with the object. If an event list was previously set for the entire file system, and a subsequent event list for an object in that file system includes an event that was set for the file system, events will continue to be generated based on the event list for the file system until such time as that event is disabled, in which case the event list for the object will come into play.

When a process generating an event is blocked awaiting a response from a DM application, the sleep mode can be interrupted.

In this release, AIX adopts a reasonably reliable model of asynchronous message delivery. The number of asynchronous messages that are not delivered is limited by the amount of available memory (real or virtual) configured on the system. If the number of messages exceeds the amount of available memory, asynchronous messages that are not delivered will be lost. JFS2 generates asynchronous namespace event messages for all corresponding operations whether they succeed or fail.

JFS2 provides interfaces that allow preallocation and direct control of metadata within a file system. Use of these interfaces with either the **MM_ALLOC** or **MM_RECORD** modes generates a DMAPI write event for the specified offset and length.

If a value is not specified for the mask to the **dm_get_bulkall**, **dm_get_bulkattr**, **dm_get_dirattrs**, and **dm_get_fileattr** subroutines (for example, it is set to zero), JFS2 will return all fields in the **dm_stat** structure.

JFS2 does not use the *respbufp* parameter of the **dm_respond_event** subroutine. If specified, the content of the buffer is undefined when the functions return.

Since JFS2 overloads the *dm_ctime* and *dm_dtime* parameters (**DM_CONFIG_DTIME_OVERLOAD** is true), the *setdtime* parameter of the **dm_set_dmattr** subroutine is ignored.

At the time a file is memory mapped (for example, when the **mmap(2)** call is executed), any nonresident portions of a file must be made resident by the DM application. To notify the application of the mapping, JFS2 will generate a read or write event corresponding to the mode and region being mapped.

mksysb

If your system fails to boot from a **mksysb** tape, you may have encountered a problem which can be identified and resolved with the following instructions. Affected systems include all CHRP architecture systems, which started with the model F50. Access the firmware command line prompt, which usually appears as an option in the SMS menus. At the firmware command line prompt, type the following two commands:

```
setenv real-base 1000000
reset-all
```

The system will then reboot, and you will be able to boot from tape, assuming that you have an otherwise valid boot image on your tape media.

mksysb Command

The way the **mksysb** command restores data through system backups that has changed.

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

If devices were removed or replaced from the system after the backup was created, their information will be restored when installing a backup, and the system will show these devices in a defined state.

These enhancements do not affect *cloning* (installing the backup onto other systems).

Known Problem with ksh and ksh93

For 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 which may result in multiple shell processes writing to the file. The shell does not detect or prevent such race conditions.

If the right-hand pattern of a regular expression contains the at (@) or vertical bar (|) characters used with their special meaning, the expression may not be properly evaluated. In this case, use **/usr/bin/ksh93**.

With **/usr/bin/ksh93** input redirection using the ″here document″ may not be successful. To work around this problem, **/usr/bin/ksh** can be used.

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

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

New Flag for the mklv Command

-T O   Uses a new device sub type of logical volume. An application (such as a database) may use this new device sub type as an indication that it may safely stop skipping the logical volume control block (lvcb) at the beginning of the logical volume (lv).

Perl

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

- **perl.rte** 5.8.0 (version 5.8)
- **perl.man.en_US**

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

```
http://www.perl.com/language/misc/Artistic.html
```

The **perl.rte** fileset will be installed automatically.

Note: IBM continues to ship Perl, but does not support it.

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/opt/perl5/man** and **/usr/opt/perl5/man64** directories.

AIX 5L Version 5.2 introduces Perl 5.8.0. 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 5.8.0.

C99 Language Interfaces

AIX 5L Version 5.2 system libraries and headers include new interfaces required by the ISO/IEC 9899:1999(E) (C99) language standard and the Single Unix Specification, Version 3. Some of the new interfaces may have the same names as symbols in existing programs. The new interfaces may be hidden by specifying `-D_NOISOC99_SOURCE` when 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.

IBM AIX Developer Kit, Java ™ 2 Technology Edition, Version 1.3.1

IBM AIX Developer Kit, Java ™ 2 Technology Edition, Version 1.3.1 is released in Java131.* filesets. For more information, see the **/usr/java131/README.HTML** file.

AIX Developer Kit, Java 2 Technology Edition, Version 1.3.1, 32-bit Version for POWER is included with the AIX base operating system. The 64-bit version is available on both the AIX 5L Version 5.2 Bonus Pack and the AIX Java Web site at http://www.ibm.com/developerworks/java/jdk/aix.

You can dynamically reconfigure a logical partition (LPAR) running a Java 1.3.1 application. Note that decreasing the number of CPUs or real memory allocated to an LPAR will likely degrade the performance of a Java application, but the application should continue to run.

The IBM 32-bit SDK for AIX, Java 2 Technology Edition, Version 1.4 and IBM 64-bit SDK for AIX, Java 2 Technology Edition, Version 1.4 are both supported on AIX 5L Version 5.2. You can download these products from the AIX Java Web site. Please install all the Java service refreshes. See the **How and where to get fixes** link from the Java 1.4 download page (click on the **Download** link on http://www.ibm.com/developerworks/java/jdk/aix. Then select the **1.4** link for the 32-bit version or the **1.4 64-bit** link for the 64-bit version). As with Java 1.3.1, you can dynamically reconfigure a LPAR running Java 1.4.

The AIX Developer Kit, Java 2 Technology Edition, Version 1.2.2 is supported on AIX 5L Version 5.2. Please install all the Java service refreshes. See the **How and where to get fixes** link from the Java 1.2.2 download page (click on the **Download** link on http://www.ibm.com/developerworks/java/jdk/aix. Then select the **1.2.2** link). You can dynamically reconfigure a LPAR running Java 1.2.2, with one exception: if the LPAR has only 1 CPU, AIX will reject any reconfiguration request that tries to increase the number of CPUs allocated to the LPAR.

AIX Developer Kit, Java Technology Edition, Version 1.1.8 and AIX Developer Kit, Java 2 Technology Edition, Version 1.3.0 are not supported on AIX 5L Version 5.2.

Location Change for License Use Management (LUM) Commands

The location of the **i4blt**, **i4cfg**, **i4target**, and **i4tv** License Use Management commands has changed from the **/var/ifor** directory to the **/usr/opt/ifor/ls/os/aix/bin** directory.

Note: To configure LUM, use the **i4cfg -script** command.

The **/usr/lib/netls/conf** directory does not exist on AIX 5L Version 5.2. The only valid directory for the nodelock file is the **/var/ifor** directory.

Performance Analysis and Tuning Tools

The following performance tuning tools are new in AIX 5L Version 5.2:
- **vmo**, **ioo**, **schedo** : The **vmtune** and **schedtune** replacement tools to set kernel tuning parameters
- **tunsave**, **tunrestore**, **tundefault**, **tuncheck** : Tunable files manipulation commands

These tools are part of the new **bos.perf.tune** fileset, which contains a set of new SMIT menus designed to manipulate AIX tunable parameters. Those menus can be reached using **smitty tuning**. Permanent tunable parameter settings are now centralized in a new stanza file called **/etc/tunable/nextboot**.

The sample code **vmtune** and **schedtune** are now shell scripts calling the **vmo**, **ioo** and **schedo** tools.

If you are migrating from a prior AIX release, read the migration information and changes introduced by the **bos.perf.tune** fileset in *AIX Kernel Tuning Parameter Modifications* in the *AIX 5L Version 5.2 Performance Management Guide*.

The following performance analysis tools are new in AIX 5L Version 5.2:
- **curt**: A thread CPU utilization reporting tool
- **splat**: A lock contention analysis tool
- **xprofiler**: A **gmon.out** GUI analyzer previously shipped with the Parallel Environment
- A set of **/proc** based tools in the new **bos.perf.proctools** fileset

JISX0213 Support for Japanese UTF-8 (JA_JP 64-bit application)

This release is a technology preview with the following restrictions:
- Range of character set for input is limited to UCS-2 code points.
- Supported character sets for display are JISX0208 and JISX0212 only.

Support requires installation of the following filesets:

bos.iconv.ucs.com Unicode Base Converters for AIX Code Sets/Fonts
bos.loc.com.JP Common Locale Support - Japanese
bos.loc.com.utf Common Locale Support - UTF-8
bos.loc.utf.JA_JP Base System Locale UTF Code Set - Japanese

For more information, refer to the **/usr/lpp/jls/doc/README.jisx0213.utf** or **/usr/lpp/jls/doc/README.jisx0213.pc** file.

CAPP/EAL4+ Technology

A CAPP/EAL4+ system, as defined in the *AIX 5L Version 5.2 Security Guide*, provides the following administrative interfaces to its security functionality:
- Identification and authentication measures (configuration of users, password settings, login configuration, etc.)
- Audit measures (configuring bin mode audition, selecting audited events, processing audit trails, etc.)
- Discretionary access control (permission bits and ACLs for file system objects, IPC mechanisms and TCP ports)
- Setting of the system time
- Running the **diag** diagnostic subsystem
- Running the **su** command to become a privileged (root) administrator

This includes the configuration files and system calls that can be used to perform the appropriate administration.

The following user interfaces to the security functionality are provided:
- The **passwd** command for changing a user's password
- The **su** command for changing a user's identity
- The **at**, **batch** and **crontab** facilities for the scheduling of command processing
- Discretionary access control (permission bits and ACLs for file system objects and IPC mechanisms)
- Login mechanisms (for example, identification and authentication mechanisms) for the system console and the supported network applications (such as, telnet, ftp)

This includes the system calls dealing with the settings of user identity or access control.

The AIX Version 5.2 CAPP/EAL4+ system runs on hardware platforms based on IBM eServer pSeries Symmetric Multiprocessor (SMP) systems using the POWER3-II CPU (p610) with 1 and 2 processors, SMP systems using the RS64 IV CPU (p660), and SMP systems using the POWER 4 CPU (p690). Peripherals that are supported are terminals and printers, hard disks and CD-ROM drives as storage devices, and streamers and floppy disk drives as backup devices. Supported network connector types are Ethernet and token ring.

Note: Do not use the **$HOME/.rhosts** configuration file for remote login and execution of commands. Administrators must inform all users of the system not to use this configuration file.

For more information on CAPP/EAL4+ Technology, see the following Web site:

`http://www.radium.ncsc.mil/tpep/library/protection_profiles/index.html`

Communications, Networking, and I/O

New Options for the no Command

tcp_finwait2	Specifies the length of time to wait in the **FIN_WAIT2** state before closing the connection, measured in half seconds. The default is 1200 half seconds (600 seconds or 10 minutes). Ranges from 0 to **USHRT_MAX**. **tcp_finwait2** is a runtime (dynamic) attribute.
tcp_inpcb_hashtab_siz	Specifies the size of the **inpcb** hash table for TCP connections.
	This table holds the **inpcb**s required for connection management and is implemented as a table of hash chains. A large table means that the linked hash chains will be small and lower the traversal time on average, but the memory footprint will be larger. Minimum value is 1, max is 999999.
	This option impacts performance and should be used with extreme caution. Please consult a performance analyst in case the value needs to be changed. The execution environment could have an influence on the value. It is strongly encouraged that you maintain the system defined defaults, because they tend to execute optimally in most environments.
	The **tcp_inpcb_hashtab_siz** option is a loadtime option that should be changed through smit or Web-based System Manager for post-AIX 5.2 systems, and takes effect after the reboot. The default size is 24999.
udp_inpcb_hashtab_siz	Specifies the size of the **inpcb** hash table for UDP connections.
	This table holds the **inpcb**s required for connection management and is implemented as a table of hash chains. A large table means that the linked hash chains will be small and lower the traversal time on average, but the memory footprint will be larger. Minimum value is 1, max is 999999.
	This option impacts performance and should be used with extreme caution. Please consult a performance analyst in case the value needs to be changed. The execution environment could have an influence on the value. It is strongly encouraged that you maintain the system defined defaults, because they tend to execute optimally in most environments.
	The **udp_inpcb_hashtab_siz** option is a loadtime option that should be changed through smit or Web-based System Manager for post-AIX 5.2 systems, and takes effect after the reboot. The default size is 24999.

7135 RAIDiant Array for AIX Requirements (after Installing AIX 5L Version 5.2)

Customers installing AIX 5L Version 5.2 on systems with attached 7135-110 or 7135-210 RAIDiant Array subsystems must also obtain and install the latest version of controller and drive microcode. This microcode update can be obtained using hardware service channels (800-IBM-SERV in the U.S.) and by requesting ECA 010. You can also order the publication for the 7135 titled, *7135 Raidiant Array for AIX: Installation Guide and Reference* (order number SC23-1742), which contains information on microcode and device driver installation.

Enable Link Polling and Time Interval for Link Polling Configuration Attributes

The IBM 10/100 Mbps Ethernet PCI Adapter (23100020) does not have an interrupt-driven mechanism to tell the device driver of any changes in its link state.

To avoid this drawback, the device driver now provides an ODM attribute to determine the status of the link. This **Enable Link Polling** attribute is disabled by default. If this attribute is enabled, the device driver will poll the adapter every time period, equal to the time in milliseconds specified in the **Time interval for Link Polling** attribute, for a change in the adapter's link state.

If the adapter's link goes down for any reason, the device driver will disable its **NDD_RUNNING** flag. When the device driver finds that the link has come back up, it will enable this **NDD_RUNNING** flag.

To work successfully, protocol layer implementations, such as EtherChannel, need notification if the link has gone down. Enable the **Enable Link Polling** flag for this adapter to accomplish this.

By enabling this flag, the performance of this adapter will decrease compared to when it is not enabled, because of the additional PIO calls that the device driver will make to the adapter every few milliseconds while it determines the link status.

Enhanced Error Handling (EEH)

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.

In order for EEH to work, the following conditions must be met:
- Must have AIX kernel support. AIX 5L for POWER Version 5.1 with the 5100-04 Recommended Maintenance package (APAR IY39794) and AIX 5L Version 5.2 with the 5200-01 Recommended Maintenance package have the latest EEH services.
- Must have AIX device driver support (dds). Most dds have full EEH recovery (with a few exceptions).
- Must have EEH capable hardware.
- Must have 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 find out if EEH will work on your system.

Supported Devices
Device Driver support for EEH is limited to the following devices that are supported by AIX 5L Version 5.2:
- Storage Adapters:
  - Fibre Channel (6227)
  - Fibre Channel (6228)
  - Fibre Channel (6239)

- – PCI SCSI RAID Devices
- – PCI Ultra SCSI Devices
- – PCI Dual Channel Ultra3 SCSI Adapter (FC 6203)
- – PCI Dual Channel Ultra2 SCSI Adapter (FC 6205)
- 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)
  - – S/390 ESCON CHANNEL PCI ADAPTER (FC 2751)
  - – IBM ARTIC960HX 4-PORT PCI ADAPTER (FC 2947)
  - – IBM ARTIC960RXD QUAD DIGITAL TRUNK ADAPTER (FC 6310)
- Encryption Adapters:
  - – IBM PCI 4758 Cryptographic Coprocessor Card (FC 4958 and 4963) *
  - – IBM Crypto Accelerator (FC 4960) *
- Graphics and Miscellaneous
  - – GXT135P Graphics Adapter (FC 2848) *
  - – USB Open Host Controller (FC 2737) *
  - – 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.

EtherChannel

With the current release, the new EtherChannel backup functionality, as well as the existing network interface backup (**netif_backup**) functionality, is supported on all Ethernet adapters.

Support for previously-unsupported Ethernet adapters is provided as fixes in the current release. The APAR numbers for these fixes are as follows:

10/100 Mbps Ethernet PCI Adapter II (1410FF01)
- APAR IY34820

10/100/1000 Base-TX PCI-X Adapter (14106902) and Gigabit Ethernet-SX PCI-X Adapter (14106802)

- APAR IY34821

10/100/1000 Base-T Ethernet PCI Adapter (14100401) and Gigabit Ethernet-SX PCI Adapter (14100401)
- APAR IY34822

Gigabit Ethernet-SX PCI Adapter and IBM 10/100/1000 Base-T Ethernet PCI Adapter

The Gigabit Ethernet-SX PCI Adapter and the IBM 10/100/1000 Base-T Ethernet PCI Adapter share the same device driver. Therefore, the device driver and diagnostic package can be obtained by installing the **devices.pci.14100401** fileset. The device driver conditionally switches code that is unique for either adapter.

Note: The **lslpp** output for devices.pci.14100401 (shown below) refers to the Gigabit Ethernet-SX PCI Adapter even though the fileset is shared by both adapters:

```
# lslpp -L | grep devices.pci.14100401
  devices.pci.14100401.diag   5.2.0.0   C   Gigabit Ethernet-SX PCI Adapter
  devices.pci.14100401.rte    5.2.0.0   C   Gigabit Ethernet-SX PCI Adapter
```

TCP Checksum Offload Feature

There is a known issue with the 10/100 Mbps Ethernet PCI Adapter II (1410ff01) when the adapter's transmit TCP checksum offload feature is enabled. Occasionally, the adapter may not compute the TCP checksum correctly for transmitted TCP/IP packets. This can cause degraded performance. In order to avoid (or remedy) this problem, the user should disable this option using the **ifconfig** command.

For example, to disable the transmit TCP checksum offload feature on interface *en0*, run the following command:

```
ifconfig en0 -checksum_offload
```

The user can also use the **ifconfig** command to check the current status of the transmit TCP checksum offload feature. For example, to check the status of this feature on interface *en0*, run the following command:

```
ifconfig en0
```

The following is returned:

```
en0: flags=7e080863,10<UP,BROADCAST,NOTRAILERS,RUNNING,SIMPLEX,MULTICAST,GROUPRT,64BIT,CHECKSUM_OFFLOAD,
CHECKSUM_SUPPORT,PSEG>
```

The **CHECKSUM_OFFLOAD** flag shown above indicates that the transmit TCP checksum offload feature is currently enabled for interface *en0*. If this feature is disabled, the **CHECKSUM_OFFLOAD** flag will not be present.

IP Security Hardware Acceleration

The 10/100 Ethernet PCI Adapter II (FC 4962) offers standards-based IP security and is designed to offload IP Security functions from the AIX base operating system. When the 10/100 Ethernet PCI Adapter II is present in the AIX system, IP Security offload must be enabled. This can be done through SMIT.

After it is enabled, the IP Security stack uses the following capabilities of the adapter:
- Encryption and decryption using DES or Triple DES algorithms
- Authentication using the MD5 or SHA-1 algorithms
- Storage of the security association information

The functions on the adapter will be used instead of the software algorithms. The IP Security hardware acceleration feature is available on level 5.1.0.25 or later of the **bos.net.ipsec.rte** and **devices.pci.1410ff01.rte** software.

IPX/SPX Protocol Support

IPX/SPX protocol support is provided in the **ipx.base** package. The **ipx.base** package is supported on the 32-bit kernel only.

Removal of Support for Devices

The following devices are not supported on AIX 5L Version 5.2 with the 5200-01 Recommended Maintenance package:

- **PCI FDDI I/O** (FC 2741, FC 2742, and FC 2743) is not supported on AIX 5L Version 5.2 with the 5200-01 Recommended Maintenance package or later.
- **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 5.2. After a migration to AIX 5L Version 5.2, or when AIX 5L Version 5.2 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 5.2.

devices.artic960.5.2.0.0

The **devices.artic960.5.2.0.0** fileset provides support for the following IBM ARTIC960 adapters:

- PCI ESCON Adapter (FC 2751)
- 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 2751, 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.

Support will be pre-loaded on factory shipped machines that ship with AIX 5L Version 5.2 with the 5200-01 Recommended Maintenance package, and on the AIX 5.2 with the 5200-01 CD.

For customers running AIX 5L Version 5.2 who desire to install a new PCI IBM ARTIC960 adapter, support will be provided on an *Additional Device Software Support* (*MES*) CD that will be shipped with the PCI IBM ARTIC960 adapter. The **bos.rte** fileset must be upgraded to 5.2.0.10.

Previous versions of the **devices.artic960** fileset need to be un-installed before installing the **devices.artic960.5.2.0.0** fileset.

Included with the **devices.artic960.5.2.0.0** 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 the **devices.artic960.5.2.0.0** fileset is installed, several pseudo-filesets will also be installed to aid in automatic creation of various ARTIC devices. Those new filesets are:

- **devices.pci.14104300.rte**, IBM ARTIC960 Device Runtime Support
- **devices.pci.14103600.rte**, IBM ARTIC960 PCI Device Runtime Support
- **devices.pci.14107800.rte**, ARTIC960 PCIHX (ARTIC960Hx) Device Runtime Support
- **devices.pci.86806019.rte**, ARTIC960 RP/RXD (ARTIC960RxD) Device Runtime Support

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 the 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 will load the necessary adapter software onto the adapter.

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

devices.pci.14108c00

The **devices.pci.14108c00** fileset provides support for SDLC and bisynchronous protocols on the IBM ARTIC960Hx 4-Port Selectable PCI Adapter, (Feature Code 2947). When combined with the installation of **devices.artic960.5.2.0.0** fileset, EEH support is now provided. 32-bit support continues to be provided for SDLC and bisynchronous communications.

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`, 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, choose **Task Selection**.
4. At the Task Selection menu, select **Display Service Hints** and press Enter.

IBM Directory (LDAP)

To access the latest IBM Directory Version 4.1.0.0 product information, go to the following Web site:

`http://www.ibm.com/software/network/directory/`

Installation and Configuration

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

`http://www.ibm.com/software/network/directory/library`

From this Web site, read the following documents:

- *IBM Directory Version 4.1.0.0 for AIX Installation and Configuration Guide*

- *IBM Directory Version 4.1.0.0 Server Readme*
- *IBM Directory Version 4.1.0.0 Client Readme*
- *IBM Directory Version 4.1.0.0 Readme Addendum*

Before starting **ldapxcfg**, verify that the following links exist as follows:

```
/usr/ldap/db2 -> /usr/lpp/db2_07_01
/usr/ldap/lib/libdb2.a -> /usr/lpp/db2_07_01/lib/libdb2.a
```

If these links are not present, create the links as follows:

```
ln -s -f /usr/lpp/db2_07_01/lib/libdb2.a /usr/ldap/lib/libdb2.a
ln -s -f /usr/lpp/db2_07_01 /usr/ldap/db2
```

Installing DB2 UDB7.2: If reinstalling, force-installing, or installing an update to the DB2 UDB 7.2 **db2_07_01.msg.ja_JP** fileset, you should first remove the **db2_07_01.msg.Ja_JP** and **db2_07_01.msg.ja_JP** filesets.

Directory Management Tool (DMT): It is strongly recommended that DMT is run when you are logged in as a user without root authority.

Administration
After installation of the **ldap.client** package create the following link:

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

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 doesn't work properly.

Multi-path I/O (MPIO)

Support for parallel SCSI SCSD disks was released in AIX 5L Version 5.2. In this release of AIX, support for MPIO Fibre Channel disks has been implemented. Some devices will configure differently than in previous releases, see "Multi-path I/O (MPIO)" on page 6 for additional information.

The following describes the major features added to the MPIO device capabilities for this release:

- Tracing the SCSI and Fibre Channel device drivers now includes traces within the MPIO FRAMEWORK. The MPIO FRAMEWORK is the interface layer between the AIX device driver and the path control module. The path control module (PCM) determines which path should be used to send I/O from the device driver to the target device.
- AIX PCM supports tracing within its interface routines. The trace hook for PCMs is 0x17B00000 HKWD_PCMKE. Vendor supplied PCMs can also use this trace hook.
- Dump support for MPIO FRAMEWORK and AIX PCM data structures are included in the component dump table (CDT). If a system crashes the MPIO FRAMEWORK and AIX PCM data structures will be included in the dump and can be viewed using the **kdb** command.
- AIX PCM health checking can test device connections (paths). This capability can re-enable failed paths. In the previous release of MPIO a user had to run the **chpath** command to re-enable failed paths.

Additional information about Multi-path I/O is available in the following AIX publications:

AIX 5L Version 5.2 System Management Concepts: Operating System and Devices in the section titled *Multi-path I/O.*

AIX 5L Version 5.2 System Management Guide: Operating System and Devices in the section titled *MPIO Devices.*

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

This release of AIX supports Dynamic Tracking and Fast I/O Failure of Fibre Channel Devices.

Dynamic Tracking allows the user to perform certain prescribed 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 bringing devices offline.

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

Support for these two features can be obtained by installing APAR IY37183. This APAR will install the **/usr/lpp/bos/README.FIBRE-CHANNEL** file. More information concerning these features can be found in this readme file.

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 5.2 Kernel Extensions and Device Support Programming Concepts* (pay special attention to the *Required FCP and iSCSI Adapter Device Driver ioctl Commands* and *Understanding the scsi_buf Structure* sections) for changes necessary to support Dynamic Tracking.

SMBFS

Acknowledgements
The SMBFS product includes software developed by Boris Popov.

Graphics

Spaceball Installation

1. Install SpaceWare as normal.
2. After SMIT starts during installation, select **enable spaceball**. The spaceball should ″beep.″
3. Select **disable spaceball**.
4. Exit SMIT.
5. Open a shell (you can use the shell that started the installation process) and type:

   ```
   /etc/SpaceWare/spaceballd -d /dev/tty0 -GL &
   ```

GL 3.2 Removed

The GL 3.2 graphics application programming interface (API) has been removed from AIX 5L Version 5.2 and is not supported on this release. Applications that use the GL 3.2 subroutines should migrate to use the OpenGL API subroutines.

System Management

CSM

You can access the CSM documentation (*Hardware Control Guide*, *Software Planning and Installation Guide*, and the *CSM for AIX Administration Guide*) by going to one of the following Web sites:

```
http://www.ibm.com/shop/publications/order
```

OR

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

Note: Before installing CSM, it is imperative that you obtain and read the *Software Planning and Installation Guide* for specific procedural steps to use when installing CSM. Installation of this product will be impossible if you do not use the documented procedures in this book.

Software Requirements
The following software is required for CSM:
* CSM FOR AIX 5L (5765-F67)

The following APARs are required for CSM on AIX 5L Version 5.2:
* IY34493
* IY34724
* IY34725

The following APAR is required for CSM on AIX 5L Version 5.1 nodes:
* IY34707

The cluster management server must be a pSeries machine running AIX 5L Version 5.2. The other machines within the cluster are referred to as managed nodes and can be pSeries machines running AIX 5L Version 5.2, or AIX 5L for POWER Version 5.1 with the 5100-03 Recommended Maintenance package.

CD/DVD Automount Facility

AIX 5L Version 5.2 provides a CD/DVD automount facility (**cdromd**) with the **bos.cdmount** fileset. To have the **cdromd** daemon enabled on each system startup, add the following line to the **/etc/inittab** file:

```
cdromd:23456789:wait:/usr/bin/startsrc -s cdromd
```

The **cdromd** daemon can interfere with scripts, applications, or instructions that attempt to mount the CD or DVD device without first checking to see if the device is already enabled. A resource or device busy error will occur in such a condition. Use the **cdumount** or **cdeject** commands to unmount the device so that you can mount the device as specified in the program or instructions. Alternatively, use the **cdcheck -m** or **mount** commands to determine the current mount point of the device.

For further information, see the **cdromd** command documentation in *AIX 5L Version 5.2 Commands Reference, Volume 2* available at the Information Center.

Notes:
1. An option has been added to the **cdutil** command to suspend the management of a device by the **automounter** daemon without causing the device to eject the media. Type `cdutil -s -k device` to suspend management without ejecting the media. By default, when the **bos.cdmount** fileset is installed, the **cdromd** entry in the **/etc/inittab** file is not enabled.
2. AIX 5L Version 5.2 system management documentation incorrectly states that CDs and DVDs are mounted by default.
3. When the **cdromd** command is enabled and an installation is performed that requires another volume, the CD will be ejected and the user will be prompted to enter the next volume. If this behavior is undesired, stop the **cdromd** command during installation. Normal behavior during multi-volume installation is to unmount the CD and ask for the next volume without ejecting.

SVR4 Printing

For help with configuring and setting up this function see the *Printing for Fun and Profit under AIX 5L* (IBM Redbook number SG24-6018-00) at the following Web site:

```
http://www.redbooks.ibm.com/
```

Using the mkprtldap Command with IBM Directory 5.1
For IBM Directory 5.1 and later, the IBM Directory (LDAP) server needs to be installed and configured on the system before the **mkprtldap** command can be run to configure System V print-specific information.

The **mkprtldap** command still supports configuring an IBM Directory 4.1 server on the system if one has not been configured before. The initial Directory configuration that is supported by the **mkprtldap** command in IBM Directory 4.1 is configuring the Administrator's DN/password and the DB2 database when they have not been configured before. This functionality is not supported by the **mkprtldap** command for IBM Directory 5.1 and later.

The **mkprtldap** command is only supported for IBM Directory 5.1 on AIX 5L Version 5.2 with the 5200-01 Recommended Maintenance package and later. On AIX 5L Version 5.2, the **mkprtldap** command is not supported with IBM Directory 5.1 and will not work.

Parallel Printer Cable Selection

The parallel printer cable must be changed to a cable that is IEEE1284-compliant if all of the following 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 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.

Web-based System Manager

Pattern Matching
Several changes have been made to the pattern matching capability of the **Find** and **Filter** actions. The AIX Version 5.1 behavior of pattern matching was to locate a substring of the pattern. The AIX Version 5.2 behavior allows the use of one or more asterisk (*****) characters to represent complex pattern matches similar to the Korn shell.

Because of this change, patterns used in AIX Version 5.1 will not work the same in AIX Version 5.2. In AIX Version 5.1, a pattern of *abc* would match property values that contained the substring *abc*. In AIX Version 5.2, a pattern of *abc* will only match property values that are exactly *abc*. In order to get the AIX Version 5.1 behavior, specify the pattern as **abc**.

In addition to the change in the pattern matching rules, the pattern matching operator in the **Filter** dialog has changed from **contains** to **matches**.

Viewing Remote Queue Status
In Web-based System Manager, the **All Print Queues** view for AIX remote printers can inaccurately indicate a problem with a remote queue. Check the actual status of the queue from the command line by typing the following command:

```
enq –q –P queue
```

If the command returns the status that the queue is ready, printing will function normally.

Inventory Scout

Inventory Scout Version 2.1 adds a new Web-based System Manager GUI function to aid customers with the maintenance of systems and device firmware. It can be enabled by installing the **invscout.websm** fileset.

Tivoli Management Agent (TMA), Version 3.7.1.0

The foundation for the Tivoli Enterprise Software product suite and third-party management products is the Tivoli Management Framework. This framework provides the graphical desktop, object-oriented databases, and base services used by other management products. TMA is a less substantial version of the Tivoli Management Framework for installation on client systems.

For more information about Tivoli Systems products and services, including how to become a Tivoli Systems customer, do one of the following:

- Call 1-800-2TIVOLI (to be connected to your local sales office)
- Visit the Tivoli Home page at the following Web address:

    `http://www.tivoli.com`

Other Software

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

`http://www.ibm.com/servers/aix/products/ibmsw/list/`

Compilers

VisualAge C++ Professional for AIX

VisualAge C++ Professional for AIX, V6.0.0, is a fully supported version of this product. Install Version 6.0.0 by using your existing Version 6.0.0 CD's, and then applying APARs IY34533, IY34534, IY34536, IY34538 and IY34623 to complete the installation on AIX 5L Version 5.2.

VisualAge C++ Professional for AIX, V5.0.2, is a fully supported version of this product. Install Version 5.0.2 by using your existing Version 5.0.2 CD's, and then applying APARs IY34533, IY34534, IY34535, IY34537 and IY34623 to complete the installation on AIX 5L Version 5.2.

C for AIX

C for AIX, V6.0.0 is a fully supported version of this product. Install Version 6.0.0 by using your existing Version 6.0.0 CDs, and then applying APARs IY34533, IY34534, IY34536 and IY34623 to complete the installation on AIX 5L Version 5.2.

C for AIX, V5.0.2 is a fully supported version of this product. Install Version 5 by using your existing Version 5.0.2 CD's, and then applying APARs IY34533, IY34534, IY34535 and IY34623 to complete the installation on AIX 5L Version 5.2.

XL Fortran for AIX

XL Fortran for AIX, Version 8.1, is a fully supported version of this product. Install Version 8.1 by using your existing Version 8.1 CDs, and then applying APARs IY34533, IY34534, IY33757, IY33758 and IY34623 to complete the installation on AIX 5L Version 5.2.

XL Fortran for AIX, Version 7.1.1 is a fully supported version of this product. Install Version 7.1.1 by using your existing Version 7.1.1 CDs, and then applying APARs IY34533, IY34534, IY33755, IY33756 and IY34623 to complete the installation on AIX 5L Version 5.2.

AIXlink/X.25

AIXlink/X.25 Version 2.0.2

AIXlink/X.25 Version 2.0.2 is supported on AIX 5L Version 5.2 with the 5200-01 Recommended Maintenance package (APAR IY29795) or later which is included on the 5/2003 or later AIX update CD.

For more information about supported adapters and about configuration and installation, see the following publications:

- *AIXlink/X.25 LPP Guide and Reference*
- *X.25 Installation and Configuration*

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

AIX Fast Connect Version 3.1.2

AIX Fast Connect documentation is available at the following Web site:

```
http://publib16.boulder.ibm.com/pseries/en_US/aixbman/fastcon/fastcontfrm.htm
```

The latest updates of the product are described in the **/etc/cifs/README** file, which is installed with AIX Fast Connect.

Quick Start
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

If you are using Communications Server for Systems Network Architecture (CS/AIX or SNA) support, use the following level:
- Communications Server for AIX, Version 6.1 (6.1.0.0 or later)

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

DCE for AIX

DCE 3.1 for AIX is not supported for the 64-bit kernel.

Removal of Support for the libipfx.a Library

AIX 5L no longer provides or supports the services provided in AIX Version 4 by the IPFX (Information Presentation Facility) LPP. In particular, the interfaces provided by the **/usr/lib/libipfx.a** IPFX library are no longer present.

It is possible that an application may still require these services. One indication is that the application fails to load with an error message that it cannot find or load the **libipfx.a** IPFX library. In this case, IBM provides an unsupported copy of the **ipfx.rte** fileset at the following Web site. This fileset is provided as a convenience only. There is no support for this fileset on AIX 5L.

```
ftp://aix.software.ibm.com/aix/efixes/ipfxrte/
```

Appendix A. AIX 5L Version 5.2 Supported and Unsupported Items

System Support

The 64-bit kernel supports 64-bit processors such as (but not limited to) the following POWER-based systems:

- RS/6000 7013 Models S70, S7A
- RS/6000 7015 Models S70, S7A
- RS/6000 7017 Models S70, S7A, S80
- RS/6000 7025 Models H80, F80
- RS/6000 7026 Models H70, H80, M80
- RS/6000 7043 Models 260, 270
- RS/6000 7044 Models 170, 270
- eServer pSeries 680 Model S85
- eServer pSeries 640 Model B80
- eServer pSeries 660 Model 6H1
- eServer pSeries 660 Model 6M1
- eServer pSeries 620 Model 6F1
- eServer pSeries 690
- eServer pSeries 610 Model 6C1
- eServer pSeries 610 Model 6E1
- eServer pSeries 670
- eServer pSeries 630 Model 6C4
- eServer pSeries 630 Model 6E4
- eServer pSeries 650 Model 6M2
- eServer pSeries 665 Model 651

AIX 5L Version 5.2 Unsupported Devices and Machines

- 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
  - 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 (feature code (FC) 2855 only)
  - – GXT800M
  - – GXT1000
  - – MVP MULTIPCI Adapter
  - – S15
  - – VIDEO OUTPUT OPTION (#3200) (FC 7254)
  - – 7250 ATTACH Adapter (FC 2856)

AIX 5L Version 5.2 Unsupported Functions and Filesets

- 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
- CPU Gard
- devices.pci.331101e0

Appendix B. 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.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:
(c) (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. (c) Copyright IBM Corp. _enter the year or years_. All rights reserved.

Trademarks

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

AIX

AIX 5L

alphaWorks

eServer

DB2

IBM

Micro Channel

PowerPC

pSeries

RS/6000

SecureWay

Tivoli

TotalStorage

TURBOWAYS

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

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

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

IBM

Printed in the USA.