

AIX 5L Version 5.3



Release Notes

AIX 5L Version 5.3



Release Notes

Note

Before using this information and the product it supports, read the information in Appendix D, "Notices," on page 67.

Ninth Edition (October 2009)

© Copyright International Business Machines Corporation 2004, 2009.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Chapter 1. Read this before installation	1
Installation tips	1
Software License Agreements (SLA)	1
Service	1
Fixes and problem-solving databases	1
Chapter 2. System requirements	3
Required hardware	3
Firmware	3
Some AIX systems might not boot from CD-ROM	3
Minimum firmware levels required for AIX 5.3	3
Firmware upgrade required to support an alternate boot device	4
I/O devices IPL limitation	6
Storage adapter microcode	7
Memory requirements	7
Host Ethernet Adapter memory requirements	7
Paging space requirements	7
Disk requirements	7
Disk capacity	8
Supported devices	8
Parallel printer cable selection	8
Supported Enhanced Error Handling (EEH) devices	9
Chapter 3. Installation, migration, upgrade, and configuration information	11
Installation	11
Installing AIX 5L Version 5.3	11
IBM Systems Director Common Agent	11
Common 32-bit kernel installation issues	12
Disk format requirement	13
Creating a bootable CD	13
Base Operating System installation options	13
rsct.opt.storagerm fileset	13
AIX Toolbox Media and NIM lpp_sources	14
Graphics software bundle requires two CDs	14
Customized bosinst.data files	14
Network Installation Management	14
CSM Server	15
IBM 32-bit SDK for AIX, Java 2 Technology Edition, Version 1.4.2	15
Perfagent Tools	16
Other installation information	16
Using mkcd and mksysb to create a 64-bit boot image	17
Troubleshooting problems with installation from mkysyb backup	17
If applications fail to run after downgrade is done	17
Migration	17
Maximum size of boot image increased	17
System migrated to AIX 5.3 might experience double boot	17
Migration from AIX 4.2.1	18
Replacements for the vmtune and schedtune commands	18
xIC.rte fileset	18
bos.clvm.enh fileset after migration to AIX 5L Version 5.3	18
KDE desktop	18
Performance monitoring API	18
SNMPv3	19

Kerberos	19
AIX Toolbox for Linux Application migration information	19
Multipath I/O (MPIO)	20
System V Printing Subsystem migration from AIX 4.3.3	24
Chapter 4. Limitations and restrictions.	25
Base Operating System (BOS)	25
console special file limitation	25
Known problems with the ksh and ksh93 commands	25
IBM Directory with Ja_JP locale	25
System management	25
Cluster Systems Management	25
Reliable Scalable Cluster Technology	26
AIX VSCSI client device driver	26
NPIV support of DS4000 storage products	26
Web-based System Manager	26
Inventory Scout, Version 2.2.0.9	27
Known limitations for POWER4 systems	28
Host Ethernet Adapter limitation	28
InfiniBand limitation	28
InfiniBand EEH limitation	29
Known problems	30
IBM 4.7 GB IDE Slimline DVD-RAM drive limitations	30
Known problem writing to DVD drive	30
Limitation on placement of boot image on hard disk	30
Machine limitations with Universal Disk Format (UDF)	31
Logical Volume Manager memory impact	31
Server Message Block File System (SMBFS) mounting	32
Restrictions using HEA under EtherChannel	32
Limitations on number of devices per vscsi client adapter	32
Other software	33
AIXlink/X.25	33
AIX Fast Connect, Version 3.2.	33
Communications Server for AIX, Version 6.1	33
Distributed Computing Environment (DCE)	33
Enterprise Identity Mapping (EIM)	34
Chapter 5. Documentation	35
AIX Information Center	35
Appendix A. AIX 5L Version 5.3 changes	37
New and Enhanced Functionality in AIX 5L Version 5.3	37
Base Operating System (BOS)	37
64-bit system identifier	37
64-bit kernel	37
32-bit kernel	37
Long user names, group names, and path names	37
System support	38
Base functionality	38
JFS2 maximum file system support	39
JFS2 file system freeze and thaw feature.	39
JFS2 sync tunables.	39
New memory allocation algorithm MALLOCTYPE=watson	40
The -l option removed from the make command	40
System performance recordings and reports	40
Reliability, availability, serviceability utilities	41

AIX Web browser Mozilla Firefox 1.5 for AIX	44
License Use Management (LUM).	45
Exclusive resource sets	46
Multiple instances of AIX on a single root volume group	46
Multiple page size support	46
USB flash drive support	47
Communications, networking, and I/O	48
IP Security	48
Asynchronous I/O fast path for CIO with JFS2	49
AIX Network Data Administration Facility	49
Internet Key-Exchange logging	49
RADIUS Server	49
Path MTU (PMTU) discovery	50
AF_INET6 sockets	50
Removal of support for devices	50
The devices.artic960 fileset	50
The devices.pci.14108c00 fileset	51
Missing resource processing	51
IBM Tivoli Directory Server (LDAP)	51
Dynamic Tracking and Fast I/O Failure of Fibre Channel devices	52
Internet Protocol (IP) over Fibre Channel	52
Sendmail, Version 8.13.4.	53
Generic Routing Encapsulation	53
AIX iSCSI software initiator	53
iSCSI Boot Support and IPv6 Support	54
Configurable IP Multipath Routing	54
Virtual SCSI client adapter	56
System management	56
AIX Network Data Administration Facility	56
Distributed Command Execution Manager (DCEM)	57
dscrctl Command	57
Enhanced nimadm command	58
Predefined XOPEN macros	58
Lightweight Infrastructure	59
Appendix B. AIX 5L Version 5.3 unsupported devices	61
Unsupported devices and machines.	61
Unsupported functions and filesets	62
Appendix C. Listing of filesets on the AIX media	65
AIX 5L for POWER Version 5.3 CD set	65
Appendix D. Notices	67
Trademarks	68

Chapter 1. Read this before installation

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

<http://www-933.ibm.com/eserver/support/fixes/fixcentral/main/pseries/aix>

These Release Notes support AIX 5L Version 5.3 with the recommended 5300-11 Technology Level.

The AIX 5L Version 5.3 Release Notes include information that helps you install AIX 5L Version 5.3. The release note information center is located at the following Web site:

Release Notes Index for AIX 5.3 and Expansion Pack

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

Installation tips

The latest installation hints and tips are available from the IBM® Subscription Service for UNIX® servers at: <http://www14.software.ibm.com/webapp/set2/subscriptions/pqvcmd>

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

Software License Agreements (SLA)

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

<http://www.ibm.com/software/sla/sladb.nsf>

Service

Fixes and problem-solving databases

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

<http://www-933.ibm.com/eserver/support/fixes/fixcentral/main/pseries/aix>

Chapter 2. System requirements

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

Required hardware

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

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

```
lscfg | grep Architecture
```

For more information about supported and unsupported items, see Appendix B, “AIX 5L Version 5.3 unsupported devices,” on page 61.

Firmware

Some AIX systems might not boot from CD-ROM

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

1. At the command prompt, type the following:

```
lscfg -vl cd*
```

2. Examine the data that is returned.

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

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

Minimum firmware levels required for AIX 5.3

The following table shows minimum system firmware levels required before installing or upgrading your system to AIX 5.3.

Model Number	Product Name	Minimum Firmware Level
7017-S70	RS/6000® Model S70	20040716
7017-S7A	RS/6000 Model S7A	20040716
7017-S80	RS/6000 Model S80	20040716
7017-S85	pSeries® 680 Model S85	20040716
7025-F50	RS/6000 Model F5	L04197
7025-F80	RS/6000 Model F80	CL050203
7025-6F0/6F1	pSeries 620 Models 6F0/6F1	CL050203
7026-B80	pSeries 640 Model B80	NAN04194
7026-H50	RS/6000 Model H50	L04197
7026-H70	RS/6000 Model H70	SST04195
7026-H80	RS/6000 Model H80	CM050203

Model Number	Product Name	Minimum Firmware Level
7026-M80	RS/6000 Model M80	MM050203
7026-6H0/6H1	pSeries 660 Models 6H0/6H1	CM050203
7026-6M1	pSeries 660 Model 6M1	MM050203
7028-6C1	pSeries 610 Model 6C1	CLT04194
7028-6C4	pSeries 630 Model 6C4	3R041029
7028-6E1	pSeries 610 Model 6E1	CLT04194
7028-6E4	pSeries 630 Model 6E4	3R041029
7029-6C3	pSeries 615 Model 6C3	3F041029
7029-6E3	pSeries 615 Model 6E3	3F041029
7038-6M2	pSeries 650 Model 6M2	3K041029
7039-651	pSeries 655 Model 651	3J041029
7040-671	pSeries 670 Model 671	3H041029
7040-681	pSeries 690 Model 681	3H041029
7043-150	RS/6000 Model 150	TCP04195
7043-260	RS/6000 Model 260	SPX04197
7043-270	RS/6000 7043 Model 270	SPH04194
7044-170	RS/6000 Model 170	SPH04194
7044-270	RS/6000 Model 270	SPH04194
7046-B50	RS/6000 Model B50	TCP04194
9112-265	IntelliStation® POWER® Model 265	CLT04194
9114-275	IntelliStation POWER Model 275	3F041029

Firmware updates are available at the following Web site:

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

If you experience problems booting your POWER4 system after updating AIX or firmware, to verify and correct the value of the real-base firmware variable, perform the following steps:

1. Go to the open firmware prompt when the system is booting by typing 8 when a beep is heard or the banner displays. If applicable, the HMC can also be used to set to boot to the SMS menu. From the SMS menu, type 0 and go to the open firmware prompt.
2. At the open firmware prompt, type:

```
printenv real-base
```

The following is a sample display:

```
----- Partition: common ----- Signature: 0x70 -----
real-base          2000000          c00000
```

The value following real-base must be c00000. If any value other than c00000 displays, run the following command:

```
setenv real-base c00000
```

3. After updating the real-base variable, reboot your system using the power switch.

Firmware upgrade required to support an alternate boot device

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

Installation

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

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

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

Note: SMS menus differ, depending on the system. On some systems, the installation device is specified under the Multiboot menus. On other systems, the installation device is specified under the Select Boot Options menus.

2. When the F5 key is selected during boot, which is the fast path to install from the CD device, *and* the CD device is not the first available device in the AIX bootlist.

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

Booting from an alternate boot device

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

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

System	Microcode Level with AIX 5.3 fix
7013 or 7015 or 7017 -S70/S7A	20040716
7017-S80/S85	20040716
7025-F50	L04197
7025 or 7026-H50	L04197
7025 or 7026-H70	SST04195
7026-B80	NAN04194
7028-6C4/6E4	3R040323 or later
7028-6C1/6E1 and 9112-265	CLT04194 or later
7029-6C3/6E3 & 9114-275	3F041021 or later
7038-6M2	3K040323 or later
7039-651	3J040528 or later
7040-671/681	3H040528 or later*
7043-150/7046-B50	TCP04195
7043-260	SPX04197
7043 or 7044-270	SPH04194
7044-170	SPH04194
9076-260	SPX04197
9076-270	SPH04194
9076-N80	NI04195
9076-N81	NI04195
9076-WCN	L04197

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

The required version of the firmware can be obtained from the following Web site:

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

I/O devices IPL limitation

Because of a firmware memory size limitation, only I/O devices in the first 144 I/O slots assigned to a logical partition or single system image partition can be used as an IPL (boot) source.

The HMC can be used to identify which devices are in the first 144 I/O slots by doing the following:

1. Select the following navigation items:
 - a. Partition Properties
 - b. Hardware tab
 - c. I/O tab
2. Click the Bus column to sort the I/O devices in ascending order.
The first 144 I/O devices in the sorted list are in the bootable adapter slots.

If you are using a partition or single system image partition with more than 144 assigned I/O slots, the following scenarios and their results and resolutions are possible.

Table 1. Scenarios for partitions with more than 144 assigned I/O slots

Scenario	Result	Resolution
Attempting to boot from a device beyond the first 144 I/O slots for installation or diagnostic purposes.	The device is not selectable as a boot source from the SMS menus.	Use a device in the first 144 I/O slots.
Booting from a device in the first 144 I/O slots, and then attempt to select a target installation device in a slot beyond first 144 I/O slots.	The boot will succeed to the installation menus, but devices beyond the first 144 I/O slots will not be listed as bootable installation targets in the AIX menus.	Select a device that is available and marked as bootable.
Using an MPIO configuration where one adapter is in the first 144 I/O slots and another adapter is in a slot beyond first 144 I/O slots. Both adapters are present at boot time.	The boot will succeed to the installation menus, and the device will be listed as bootable in AIX installation menus. The installation will proceed, but it will fail with the bootlist command failure "unable to set bootpath for all paths."	Use a device in the first 144 I/O slots for all paths.
Using DLPAR to add an adapter in a slot beyond the first 144 I/O slots, and then attempting to run the alt_disk_install command for the newly added device.	The device will not be listed as bootable.	Use a device in the first 144 I/O slots.
Using DLPAR to add an adapter in a slot beyond the first 144 I/O slots, and using the bootlist command to add the device as a bootable device (for example, by dynamically adding a redundant path to a current boot device or setting up for a network boot). Then removing the original adapter and rebooting.	The bootlist command succeeds, but the boot fails from the specified device, and AIX will not receive control.	Use a device in the first 144 I/O slots.

Table 1. Scenarios for partitions with more than 144 assigned I/O slots (continued)

Scenario	Result	Resolution
Using DLPAR to add an adapter whose probe order will make it displace a current bootable device, and then rebooting.	The boot fails, and AIX will not receive control.	Move the boot device to one of the first 144 I/O slots or remove the previously added device.

Storage adapter microcode

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

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

All these adapters support concurrent microcode download. Check for the latest adapter microcode updates at the following Web site:

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

Memory requirements

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

A general rule for a minimum current memory requirement for AIX 5L Version 5.3 is 256 MB–512 MB.

AIX 5L Version 5.3 requires the minimum current memory requirement to increase as the maximum memory configuration or the number of devices scales upward, or both. Larger maximum memory configurations or additional devices scale up the minimum current memory requirement. If the minimum memory requirement is not increased along with the maximum memory configuration, the partition hangs during the initial program load (IPL).

Host Ethernet Adapter memory requirements

Configurations containing a Host Ethernet Adapter (HEA) require more memory than the 512 MB minimum. Each logical HEA port that is configured requires an additional 102 MB of memory. The minimum memory requirement for configurations with one or more HEA ports configured, where n is the number of HEA ports, is 512MB + n*102MB.

Paging space requirements

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

Disk requirements

AIX 5L Version 5.3 with the recommended 5300-11 Technology Level requires a minimum of 4 GB of physical disk space for a default installation, which includes all devices and the Graphics bundle.

Notes:

1. The following measurements provide information about disk usage when you install AIX 5L Version 5300–11.
2. The boot logical volume is required to be 24 MB with AIX 5L with 5300-11. The pre-migration script in the **usr/lpp/bos** directory on your media or in your 5300–11 NIM SPOT will verify that you have adequate room. During the migration, hd5 is increased if necessary. The logical partitions must be contiguous and within the first 4 GB of the disk.

Disk requirements for AIX 5L Version 5.3 with the recommended 5300-11 Technology Level	
Location	Allocated (Used)
/	72 MB (62 MB)
/usr	1700 MB (1579 MB)
/var	360 MB (235 MB)
/tmp	128 MB (4 MB)
/opt	364 MB (198 MB)

Note: Before migrating, ensure that you have an adequate number of available partitions and increase the **/tmp** directory to 44 MB.

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

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

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

Disk capacity

For the AIX 32-bit kernel, the maximum supported capacity of a disk is 1 TB.

For the AIX 64-bit kernel, SAS RAID controller arrays (hdisk) are supported that have capacities that exceed 2 TB (up to the RAID array capacity limitation of the SAS RAID controller).

For additional information about SAS RAID controllers, see the SAS RAID controller for AIX Web page.

For disks other than SAS RAID controller arrays, the maximum supported capacity of a disk is 2 TB.

For information about AIX capacity limitations for logical storage, see the Limitations for logical storage management Web page.

Supported devices

Parallel printer cable selection

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

- Your system was manufactured after 1998.
- The printer is "parallel attached."
- The attached printer is not a dot-matrix printer.

- The output of the **lsdev -C -l ppa0** command contains the word IEEE1284.

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

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

Supported Enhanced Error Handling (EEH) devices

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

For EEH to work, your system must have:

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

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

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

Installation

This section contains information about installing AIX 5.3 that supplements the information contained in *AIX 5L Version 5.3 Installation and migration*. The installation guide is available online in the AIX Information Center and in printed hardcopy (SC23-4887-04).

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

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

<http://www14.software.ibm.com/webapp/set2/subscriptions/pqvcmj>

Installing AIX 5L Version 5.3

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

- Complete overwrite installation
- Preservation installation
- Migration installation

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

If your system has AIX 5.3 with 5300-00 through 5300-10 installed, you can use either the Update CD (recommended) or the base media to update to AIX 5L with 5300-11. To verify the level, run the **oslevel -r** command. To perform the update, use the **smitty update_all** command.

Attention: Using the Update CD is recommended for updates. If you use the base media to update, some ODM settings (such as SRC subsystems settings) might be lost.

Notes:

1. Because only the base install images are on the media, if you use the product media to update to AIX 5L with 5300-11, you cannot reject the software and return to the previous level.
2. The size of the boot logical volume has increased from 16 MB to 20 MB. If your current boot logical volume is 16 MB or less, the installation process tries to increase it. However, partitions in a boot logical volume must be contiguous, and within the first 4 GB on the disk. If your system does not have free space that meets these requirements, a message will indicate there is not enough space to expand **hd5** (the boot logical volume).

To install AIX 5L Version 5.3, boot your system from the product media, and follow the instructions in the *Installation and migration* in the AIX Information Center. Booting from DVD media is supported only on 64-bit systems.

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

IBM Systems Director Common Agent

Beginning with AIX 5.3 with the 5300-10 Technology Level, when you do an overwrite or a preservation operating-system installation and you use the default options (which include the Graphics bundle), the IBM Systems Director Common Agent 6.1.0.3 is installed on your system.

When the AIX system is rebooted, the Director agent and the prerequisite processes for it, like the Pegasus CIM server, are automatically enabled. If you do not want to install or enable the Director agent, see the following information for the steps needed to stop, disable, and uninstall the Director agent.

To stop the Director agent, run the following commands:

```
/opt/ibm/director/agent/runtime/agent/bin/endpoint.sh stop
stopsrc -s platform_agent
```

To start the Director agent, run the following commands:

```
startsrc -s platform_agent
/opt/ibm/director/agent/runtime/agent/bin/endpoint.sh start
```

To permanently disable the Director agent, do the following:

1. Comment out the following lines in **/etc/inittab** file:

```
platform_agent:2:once:/usr/bin/startsrc -s platform_agent >/dev/null 2>&1
```

Note: Comment out the following line only if you do not need the **cimserver** command to run other applications.

```
cimservices:2:once:/usr/bin/startsrc -s cimsys >/dev/null 2>&1
```

2. Run the following command:

```
/opt/ibm/director/agent/runtime/agent/bin/endpoint.sh stop
```

3. Run the following command:

```
/opt/ibm/director/agent/runtime/nonstop/bin/installnonstop.sh -uninstallservice
```

To permanently uninstall the Director agent components and leave the filesets that might be used by other applications, run the following command:

```
/opt/ibm/director/bin/diruninstall
```

To uninstall the filesets that remain after you run the **diruninstall** command (if they are not required for other applications), do the following:

- To uninstall the CAS agent, run the following command:

```
installp -u cas.agent
```

- To uninstall Tivoli® Guid, run the following command:

```
installp -u tivoli.tivguid
```

- To uninstall SMIS CIM providers, run the following command:

```
installp -u sysmgmt.cim.smisproviders*
```

- To uninstall the remaining CIM providers, run the following command:

```
installp -u sysmgmt.cim.providers*
```

- To uninstall the Pegasus CIM server, run the following command:

```
installp -u sysmgmt.cimserver.pegasus
```

For more information about IBM Systems Director 6.1, see the IBM Systems Director V6.1 Web page.

For additional Power systems information about IBM Systems Director 6.1, see the IBM Systems Director Best Practices Redwiki Web page.

Common 32-bit kernel installation issues

The following types of installation failures can occur during a network installation with a 32-bit network boot image:

- The client will fail to boot in a logical partition that has more than 512 GB of memory.

- When an installation is to a root volume group (rootvg) that is larger than 1.2 TB, the installation fails and the following message is displayed:

ATTENTION: The total size of a disk selected for rootvg is larger than what the running kernel supports. Select a smaller disk size or reboot using media that supports running the 64-bit kernel.

The installation failures described in this section do not occur when a 64-bit network boot image is used.

Disk format requirement

You cannot install AIX on an improperly formatted SCSI disk. AIX requires the disk to be formatted to a sector size supported by the attached SCSI controller. All AIX SCSI controllers support 512 byte sector SCSI disks. The 522 byte sector SCSI disks are only supported when they are attached to SCSI RAID controllers. If the disk has been formatted for SCSI RAID, but is not attached to a SCSI RAID controller, the disk might not configure. If the disk does configure, it might be unreadable in the AIX environment. In some instances, the certify function and the format function in AIX diagnostics can be used to reformat the disk for the attached SCSI controller.

Creating a bootable CD

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

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

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

OR

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

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

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

Base Operating System installation options

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

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

SCSI disks

Disks may be grouped by adapter or SCSI bus

IBM TotalStorage® DS4000®

Disks are grouped by disk array controller (DAC)

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

rsct.opt.storagerm fileset

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

AIX Toolbox Media and NIM lpp_sources

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

Graphics software bundle requires two CDs

Due to space constraints on the AIX base product media CDs, the graphics software bundle is now included on the *Volume 2* and *Volume 3* CDs. For installations using CD media, you are prompted for the *Volume 2* and *Volume 3* CDs if you install with the defaults for an overwrite or preservation installation (Graphics Software = yes).

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

1. Type the **smitty nim_update_add** command. The **Add Software to an lpp_source** menu is displayed.
2. Select the **Graphics** bundle for **INSTALLP BUNDLE containing packages to add**. Change the answer for Process multiple volumes? to yes, if necessary

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

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

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

Customized bosinst.data files

Beginning with the 5300–08 Technology Level, new fields, and fields that now support default values, allow the **bosinst.data** files to be used more generically.

If you are using customized **bosinst.data** files on any supported media or as NIM bosinst_data resources, refer to the bosinst.data stanza descriptions and update your files. On a current system, refer to the **/usr/lpp/bosinst/bosinst.template.README** file.

Note: Systems managed by HMCs may have the console changed from **/dev/tty0** in AIX 5.2 to **/dev/vty0** in AIX 5.3. Using **/dev/tty0** to install AIX 5.3 on the system may result in a hang. To avoid this problem, use the Default value.

Network Installation Management

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

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

For NIM **rte** installs, when using the **New and Complete Overwrite** installation method, if you are using an `lpp_source` created from a previous version of AIX 5.3 base media, and updates from the current version of AIX (from the *Update CD*), you should specify an `image_data` resource. After your NIM SPOT is updated from your `lpp_source`, you can copy the `<SPOT_location>/lpp/bosinst/image.template` to a new location, and make it your `image_data` resource. This is because the number of partitions required for **hd5** (boot logical volume) increased in AIX 5.3 with the 5300-07 Technology Level.

CSM Server

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

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

- tcl
- tk
- expect
- conserver

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

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

IBM 32-bit SDK for AIX, Java 2 Technology Edition, Version 1.4.2

IBM 32-bit SDK for AIX, Java™ 2 Technology Edition, Version 1.4.2 is released in Java14.* filesets, and it is shipped with the base operating system. For more information about this version, see the </usr/java14/docs/sdkguide.aix32.htm> file.

The 64-bit version is available on both the AIX 5L Version 5.3 Expansion Pack and the following AIX Java Web site:

<http://www.ibm.com/developerworks/java/jdk/aix/>

You can dynamically reconfigure a logical partition (LPAR) that is running a Java 1.4.2 application.

Note: If the number of processors or the amount of real memory that is allocated to an LPAR is decreased, the performance of a Java application is reduced, but the application continues to run.

The following versions are not supported:

- IBM AIX Developer Kit, Java 2 Technology Edition, Version 1.3.1, 32-bit version for POWER
- IBM AIX Developer Kit, Java 2 Technology Edition, Version 1.3.1, 64-bit version for POWER

Perfagent Tools

The **perfagent.tools** fileset is automatically installed with the base operating system installation. Four files have been moved to the new **bos.perf.fdpr** fileset that is not automatically installed. The moved files are:

- **/usr/lib/perf/fdprpro**
- **/usr/lib/perf/libfdprinst.a**
- **/usr/bin/fdpr**
- **/usr/lib/perf/libfdprinst64.a**

If you require these files, the new fileset is on the *Volume 3* CD.

Other installation information

Installation packaging formats

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

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

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

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

```
/mount_point/installp/ppc
```

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

```
/mount_point/RPMS/ppc
```

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

```
/mount_point/ismp/ppc
```

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

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

Using mkcd and mksysb to create a 64-bit boot image

When a bootable backup of a root volume group is created, the boot image reflects the currently running kernel. If the current kernel is the 64-bit kernel, the backup's boot image is also 64 bit, and it only boots 64-bit systems. If the current kernel is a 32-bit kernel, the backup's boot image is 32 bit, and it can boot both 32-bit and 64-bit systems.

Troubleshooting problems with installation from mksysb backup

Check that you have sufficient free blocks in the file systems to write temporary files. Check that each file system has at least 500 blocks free when the **mksysb** backup image is made. The system needs work space in each file system when installing from a **mksysb** backup image.

Note: Depending on the type of data or files in the file system, you might need additional blocks free. For example, if the file system has a lot of small files, an extra 4 KB is automatically allocated to allow for metadata expansion.

If applications fail to run after downgrade is done

Upgrading to a new technology level and then downgrading again to the previous level because of some issues with the new level might cause some applications to not work. This is due to the applications still using symbols from the new technology level libraries. If applications fail to run after a downgrade is done, you need to relink the applications again. Because the relink procedure is different for different applications, you should contact the appropriate application support.

Migration

This section contains information about installing AIX 5.3 that supplements the information contained in *AIX 5L Version 5.3 Installation and migration*. The installation guide is available online in the AIX Information Center and in printed hardcopy (SC23-4887-04).

Note: Installing any new level of AIX requires more disk space than previous levels. Read the Disk Requirements in Chapter 2 for the disk requirements for a basic installation and ensure that you have free space either already existing in the systems, or free partitions in the rootvg. Migrating will require slightly more free space than a basic installation.

Maximum size of boot image increased

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

System migrated to AIX 5.3 might experience double boot

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

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

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

Migration from AIX 4.2.1

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

Replacements for the vmtune and schedtune commands

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

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

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

x1C.rte fileset

If you are migrating to AIX 5.3 from AIX 4.2.x or AIX 4.3.x, check the level of the **x1C.rte** fileset by typing the following command:

```
ls1pp -L x1C.rte
```

If the **x1C.rte** level is earlier than 5.0.2.x, you must apply APAR IY17981 before migrating to AIX 5.3. Without APAR IY17981 installed, the migrated system might fail to boot.

APAR IY17981 is available from the following Web site:

<http://www-933.ibm.com/eserver/support/fixes/fixcentral/main/pseries/aix>

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

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

KDE desktop

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

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

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

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

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

Performance monitoring API

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

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

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

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

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

```
lslpp -l pmtoolkit
```

- If you get the following output:

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

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

SNMPv3

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

Kerberos

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

If you are migrating to AIX 5.3, using CD or DVD media, and have Kerberos Version 5 installed, the installation scripts will prompt you to install the **krb5.client.rte** fileset from the Expansion Pack. If you are migrating your operating system using NIM resources, add **krb5** to your **lpp_source**. The secure remote commands support Kerberos clients and servers from both Native Kerberos 5 and DCE.

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

AIX Toolbox for Linux Application migration information

If you previously installed the AIX Toolbox for Linux Applications and the level of the **rpm.rte** fileset is lower than 3.0.5.20, remove that software from the system before migrating to AIX 5.3. The Toolbox software installed with **rpm.rte** levels prior to 3.0.5.20 are incompatible with software from the AIX Toolbox installed on AIX 5L Version 5.3 because of shared library restructuring.

Remove the software if you are performing a preservation installation and you established an **/opt/freeware** file system for the Toolbox software. The files in that file system will not be automatically

overwritten during a preservation installation. To remove your existing rpm filesets, use the **destroyRPMS** tool available in the **/contrib** directory on the *AIX Toolbox for Linux Applications* CD by typing the following:

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

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

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

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

Multipath I/O (MPIO)

In AIX Version 5.3, the DS4K products are configured as FCPARRAY devices by default.

A DS4K product using the FCPARRAY driver may be migrated to the MPIO driver by using the following instructions, or all DS4K devices may be migrated by uninstalling the **devices.fcp.disk.array.rte** package and then running the **cfgmgr** command or rebooting.

The DS3K products are always configured as MPIO devices. There is no FCPARRAY support for these devices. The MPIO support for these devices removes the limitation of connecting only one adapter to each DS3/4K controller port.

Instructions to convert DS4000 disks from FCPARRAY to MPIO

The MPIO driver has been enhanced to support most DS4000 models. The MPIO driver allows more than two paths to be configured to the device, and this simplifies zoning and improves performance. For the DS4000 to operate correctly in an MPIO configuration, the following steps must be performed.

The following AIX releases and minimum service packs can be used to convert DS4000 disks from FCPARRAY to MPIO:

- AIX 5.2 with the 5200-10 Technology Level and the 5200-10-04 Service Pack
- AIX 5.3 with the 5300-06 Technology Level and the 5300-06-05 Service Pack
- AIX 5.3 with the 5300-07 Technology Level and the 5300-07-02 Service Pack
- AIX 6.1 with the 6100-00 Technology Level and the 6100-00-02 Service Pack

The following DS4000 models are supported by AIX MPIO:

- 1722-600 = DS4300 = FAStT 600
- 1742-900 = DS4500 = FAStT 900
- 1814 = DS4700/DS4200
- 1815 = DS4800

To convert DS4000 disks from FCPARRAY to MPIO, do the following:

1. Change the NVSRAM settings on DS4000 storage used in the SAN (or direct attach) by doing the following:

Note: If this step was previously performed, go to step 2. However, if DS4000 firmware was updated, this step must be performed again.

- a. From the main Enterprise Management screen, highlight the subsystem you want to change.
- b. From the Tools menu, select Execute Script to open a Script Editor window.
- c. Enter the following:


```
set controller [a] HostNVS RAMBYTE [0x06,0x27] = 0;
set controller [b] HostNVS RAMBYTE [0x06,0x27] = 0;
```
- d. From the Tools menu, use the Script Editor to select Verify & Execute.
- e. Wait for the script to complete.

Attention: Do not cancel the script operation.

When the script completes, exit the Storage Manager.

2. Install the MPIO_DS3_4K APAR/iFIX that is appropriate for the AIX release on the server, and then reboot the AIX server.
3. Determine the DS4000 model or models that you are using on your AIX host by running the following commands, where X is the hdisk number for the DS4000 Array Disk type:

```
lsdev -Cc disk | grep DS4K
lscfg -vI hdiskX
```

For example:

```
# lsdev -Cc disk
hdisk2 Available 0C-08-02 DS4K Disk Array Device

# lscfg -vI hdisk2
hdisk2 U787A.001.DPM0WY4-P1-C5-T1-W200800A0B81326B7L3000000000000 1742-900 (900) Disk Array Device

Manufacturer.....IBM
Machine Type and Model.....1742-900
ROS Level and ID.....30393134
Serial Number.....
Device Specific.(Z0).....0000053245004032
Device Specific.(Z1).....
```

4. Use the following table and the Machine Type and Model output from step 3 to determine your DS4000 Array Disk type.

Table 2. DS4000 Array Disk Cross Reference

Machine Type-Model	Array Disk	Storage Server
1722-600	DS4300	FAST 600
1742-900	DS4500	FAST 900
1814	DS4700/DS4200	
1815	DS4800	

5. Check the AIX systems driver support and state by running the following command:

```
# 'manage_disk_drivers'
```

The output from this command displays the DS4000 devices that are supported by MPIO and RDAC, and indicates if they are currently configured for MPIO or RDAC enablement.

For example:

```
# manage_disk_drivers
1: DS4300: currently RDAC/fcpararray; supported: RDAC/fcpararray, MPIO
2: DS4500: currently RDAC/fcpararray; supported: RDAC/fcpararray, MPIO
3: DS4700/DS4200: currently RDAC/fcpararray; supported: RDAC/fcpararray, MPIO
4: DS4800: currently RDAC/fcpararray; supported: RDAC/fcpararray, MPIO
```

Note: If the AIX system has more than one type of DS4XXX models attached, run the **manage_disk_drivers** command for each model type.

In the example in step 3 on page 21, hdisk2 is a DS4500. In the example in step 5 on page 21, DS4500 is currently enabled for RDAC/fcparray.

To change the driver that manages the DS4000, run the following commands in order, where *X* in the first command is the number of the array disk stanza from the output of the **manage_disk_drivers** command. For example: for DS4700, enter # `manage_disk_drivers - c 3`; for DS4800, enter # `manage_disk_drivers - c 4`, etc.

```
# manage_disk_drivers - c X
```

To verify the change, answer *Y* when prompted.

```
# bosboot -a
# shutdown -Fr
```

6. After the server reboots, all of the hdisk devices that were converted will be configured under the MPIO driver. All of the hdisks that are connected to DS4000 should be displayed as MPIO Other DS4K Array Device as indicated by the **lsdev -Cc disk** output.
7. To verify that hdisks are mapped to the appropriate LUNs on DS4000, run the **mpio_get_config -A** command. This command is similar to the **fget_config** command for the RDAC driver. The **mpio_get_config** command is for informational purposes only; you cannot change the LUN ownership using this command.

Instructions to convert DS4000 disk driver from MPIO to FCPARRAY

The following AIX releases and minimum service packs can be used to convert the DS4000 driver from MPIO to FCPARRAY:

- AIX 5.2 with the 5200-10 Technology Level and the 5200-10-04 Service Pack
- AIX 5.3 with the 5300-06 Technology Level and the 5300-06-05 Service Pack
- AIX 5.3 with the 5300-07 Technology Level and the 5300-07-02 Service Pack
- AIX 6.1 with the 6100-00 Technology Level and the 6100-00-02 Service Pack

To convert the DS4000 driver from MPIO to FCPARRAY, do the following:

1. Install the MPIO_DS3_4K APAR/iFIX that is appropriate for the AIX release on the server.
2. Change the NVSRAM settings on DS4000 storage used in the SAN (or direct attach) by doing the following:

Note: If this step was previously performed, go to step 3. However, if DS4000 firmware was updated, this step must be performed again.

- a. From the main Enterprise Management screen, highlight the subsystem you want to change.
- b. From the Tools menu, select Execute Script to open a Script Editor window.
- c. Enter the following:

```
set controller [a] HostNVSRAMBYTE [0x06,0x27] = 0;
set controller [b] HostNVSRAMBYTE [0x06,0x27] = 0;
```

- d. From the Tools menu, use the Script Editor to select Verify & Execute.
- e. Wait for the script to complete.

Attention: Do not cancel the script operation.

When the script completes, exit the Storage Manager.

3. To migrate from MPIO to RDAC, there must be no more than two connections between the server and the storage (one per controller). Verify that this requirement is met by verifying that there is a separate zone for each HBA and controller pair. Refer to the RDAC documentation for valid configuration details.
4. After the installation completes successfully, reboot the system.
5. Determine the DS4000 model or models that you are using on your AIX host by running the following command where *X* is the hdisk number for the DS4000 Array Disk type:

```
lsdev -Cc disk | grep DS4K
lscfg -v1 hdiskX
```

For example:


```
# lsdev -Cc disk
hdisk2 Available 0C-08-02 DS4K Disk Array Device

# lscfg -v1 hdisk2
hdisk2 U787A.001.DPM0WY4-P1-C5-T1-W200800A0B81326B7L3000000000000 1742-900 (900) Disk Array Device

Manufacturer.....IBM
Machine Type and Model.....1742-900
ROS Level and ID.....30393134
Serial Number.....
Device Specific.(Z0).....0000053245004032
Device Specific.(Z1).....
```

- Use the following table and the Machine Type and Model output from step 5 on page 22 to determine your DS4000 Array Disk type.

Table 3. DS4000 Array Disk Cross Reference

Machine Type-Model	Array Disk	Storage Server
1722-600	DS4300	FAStT 600
1742-900	DS4500	FAStT 900
1814	DS4700/DS4200	
1815	DS4800	

- Check the AIX systems driver support and state by running the following command:

```
# 'manage_disk_drivers'
```

The output from this command displays the DS4000 devices that are supported by MPIO and RDAC, and indicates if they are currently configured for MPIO or RDAC enablement.

For example:

```
# manage_disk_drivers
1: DS4300: currently RDAC/fcpararray; supported: RDAC/fcpararray, MPIO
2: DS4500: currently RDAC/fcpararray; supported: RDAC/fcpararray, MPIO
3: DS4700/DS4200: currently RDAC/fcpararray; supported: RDAC/fcpararray, MPIO
4: DS4800: currently RDAC/fcpararray; supported: RDAC/fcpararray, MPIO
```

Note: If the AIX system has more than one type of DS4XXX models attached, run the **manage_disk_drivers** command for each model type.

In the example in step 5 on page 22, hdisk2 is a DS4500. In the example in step 7, DS4500 is currently enabled for RDAC/fcpararray.

To change the driver that manages the DS4000, run the following commands in order, where X in the first command is the number of the array disk stanza from the output of the **manage_disk_drivers** command. For example: for DS4700, you would enter # **manage_disk_drivers - c 3**; for DS4800, you would enter # **manage_disk_drivers - c 4**, etc.

```
# manage_disk_drivers - c X
```

To verify the change, answer Y when prompted

```
# bosboot -a
# shutdown -Fr
```

- After the server reboots, all of the hdisk devices that were converted will be configured under the RDAC driver.

IBM Blade Center Issues

To ensure that the storage devices are properly configured on the IBM Blade Center, change the **init_link** setting on the Fibre Channel adapter from arbitrated loop to point to point.

Enter the following commands to change the **init_link** setting on the Fibre Channel adapter:

```
chdev -l fcsX -ainit_link=pt2pt -P
bosboot -aD
cfgmgr
```

System V Printing Subsystem migration from AIX 4.3.3

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

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

Chapter 4. Limitations and restrictions

Base Operating System (BOS)

console special file limitation

The EIO errors are restricted in some cases when the connection to the console is lost. The EIO errors that result from the console write function are restricted because the streams file is marked when a tty that is connected to a modem is disconnected. The console redirects the writes to another console, and any writes to the file descriptor are returned with EIO errors. These EIO errors continue until a process reopens the console.

Known problems with the ksh and ksh93 commands

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

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

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

IBM Directory with Ja_JP locale

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

System management

Cluster Systems Management

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

<http://www-03.ibm.com/systems/clusters/support/index.html>

Click the **Cluster info center** link, and then click the **Cluster Systems Management** link in the **Overview** section.

Software requirements

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

Reliable Scalable Cluster Technology

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

Restriction for Japanese locales

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

Service Resource Manager (ServiceRM)

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

AIX VSCSI client device driver

AIX has added support for Virtual SCSI fast fail over. This function requires VIOS level 2.1 or later, and is enabled by changing the VSCSI **fast_fail** attribute to **true**. The change occurs on the next reboot.

NPIV support of DS4000 storage products

Some models of the DS4000 product family can be configured using the FCPARRAY or MPIO driver:

- If the device is configured using the MPIO driver, there are no restrictions.
- If the device is configured using the FCPARRAY driver, then the following restriction applies:
 - If the device has the reservation policy attribute set to any value other than **no_reserve**, the LPARs that the device is attached to will fail LPAR migration operations.

Web-based System Manager

Remote client management

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

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

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

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

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

<http://www.ibm.com/software/webservers/httpservers/>

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

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

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

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

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

Inventory Scout, Version 2.2.0.9

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

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

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

The new microcode management feature does the following:

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

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

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

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

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

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

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

<http://www-933.ibm.com/eserver/support/fixes/fixcentral/main/pseries/aix>

Known limitations for POWER4 systems

Adapters

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

CPU Gard

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

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

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

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

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

System p[®] 690 Memory

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

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

Host Ethernet Adapter limitation

AIX Dynamic Logical Partitioning (DLPAR) of Integrated Virtual Ethernet (IVE), also known as Host Ethernet Adapter (HEA), is not supported in AIX 5.3 with the 5300-06 Technology Level and AIX 5.2 with the 5200-10 Technology Level. Support for this feature is enabled in AIX 5.3 with the 5300-07 Technology Level and AIX 6.1 with the 6100-00 Technology Level and later releases.

InfiniBand limitation

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

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

- IbDebugChk - returns the current debug status
- IbDebugOn - enables the debug tracing

- `IbDebugOff` - disables the debug tracing

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

ibstat command

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

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

- d** Displays current debug setting
- h** Displays ibstat command usage
- i** Displays Network Interface Information
- n** Displays IB node information
- p** Displays IB port information
- v** Displays all IB device information

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

Device Name

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

Port State

Displays the current state of each HCAD port

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

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

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

InfiniBand EEH limitation

The 4X IB Dual Port DDR Exp CD, CFF/PCIE adapter (feature 8258) does not support full EEH recovery. If you get a bus error (`MXIB_EEH_HARDWARE_ERROR` in the error log), you might have to manually perform the recovery steps. These steps may include manually reconfiguring the adapter or replacing the adapter. To manually reconfigure the adapter, do the following:

1. Unconfigure the adapter by using the **rmdev** command (for example, `rmdev -l iba0`).
2. Configure the adapter by using the **mkdev** command.

Notes:

1. If the adapter instance is busy (a program or kernel extension has called `open` on the adapter), you might not be able to unconfigure the adapter. If you cannot unconfigure the adapter, reboot the system.
2. If the bus error (`MXIB_EEH_HARDWARE_ERROR`) is present after a successful manual recovery, replace the adapter.

Known problems

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

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

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

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

IBM 4.7 GB IDE Slimline DVD-RAM drive limitations

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

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

Known problem writing to DVD drive

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

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

You can verify the level of cdrecord by running the `ls1pp -L cdrecord` command. It needs to be at least at level 1.9-6, which is available on this current AIX release or you can get it from:

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

Limitation on placement of boot image on hard disk

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

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

Firmware on the affected machines cannot read the boot image from the hard disk if any part of the boot image is located beyond the 4 GB boundary on the hard disk. This is not a problem for most customers because the AIX installation process creates the boot logical volume at the beginning of the disk. This is achieved by using the **-a** flag with the **mklv** command and specifying **e** (which corresponds to **edge**) as the parameter for the **-a** flag. Using the **mklv** command with this parameter results in the boot logical volume being created at the edge of the hard disk, and the resulting address that the firmware uses to read the boot image will be within a safe range. The AIX installation process has always created the boot logical volume near the edge of the hard disk because that region of the hard disk has the slowest access time, and this allows other regions of the hard disk to be used by file systems that can benefit from increased performance.

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

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

An example of this calculation follows:

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

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

Machine limitations with Universal Disk Format (UDF)

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

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

Logical Volume Manager memory impact

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

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

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

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

With AIX 5.3 Technology Level 5300-05, the extra LVM memory footprint related to memory affinity only applies to hardware platforms that support memory affinity. Thus, disabling the `memory_affinity` tunable on a system, which doesn't support memory affinity, will not have any effect on the LVM memory footprint. On systems where memory affinity is unsupported, the LVM only requires approximately 270 KB of additional pinned memory per physical volume on AIX 5.3 Technology Level 5300-05 than on AIX 5.2 or AIX 5.1.

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

Server Message Block File System (SMBFS) mounting

On the mounted file system, a copy operation of one file to another is successful for a 4 GB + 4096 bytes or less file size. A warning message is printed for files greater than 4 GB + 4096 bytes, and the maximum 4 GB + 4096 original files are copied to the destination.

Restrictions using HEA under EtherChannel

Host Ethernet Adapter logical ports are only supported under EtherChannel if all adapters within the EtherChannel are HEA logical ports (including the backup adapter, if any). Consequently, having HEA logical port adapters intermixed with physical Ethernet adapters or Virtual I/O Ethernet Adapters in the same EtherChannel is not supported.

When using multiple HEA logical ports as primary adapters in an EtherChannel, the physical ports associated with the HEA logical ports must also be placed in an EtherChannel in the Ethernet switch. Consequently, all partitions that use HEA logical ports going to the same HEA physical ports must also be placed in an EtherChannel.

For example, assume that Partition 1 is configured as follows:

- A logical HEA port out of physical HEA port 0
- A logical HEA port out of physical HEA port 1
- An EtherChannel created using the logical HEA ports listed above

If another partition on the same system that needs to use a logical HEA port out of physical HEA port 0 or out of physical HEA port 1, you must create an EtherChannel for the partition over both of the logical HEA ports, similar to the configuration of Partition 1. Attempting to use either of those logical HEA ports as stand-alone in other partitions might cause connectivity problems, because packets might not be delivered to the correct logical HEA port.

The aforementioned configuration restriction does not exist for using logical HEA ports in a Network Interface Backup configuration (1 primary and 1 backup), since the physical HEA ports do not require specific configuration on the Ethernet switch.

Limitations on number of devices per vscsi client adapter

Prior to AIX Version 5.3 with the Technology Level 5300-07, the limit on the number of open devices (disks and optical devices) per vscsi client adapter instance was 169. Because there are problems with performance with this number of open devices, the algorithm was changed.

The vscsi adapter now performs as follows:

- When the number of open devices per adapter reaches 85, and the Virtual I/O Server has been configured to have enough resources to support it, the adapter attempts to allocate more memory when a child device is opened.
- If the memory allocation is successful, the adapter allows up to 255 devices to be open at one time. If the memory allocation is not successful, the adapter allows up to 127 devices to be open at one time.
- If the additional memory is obtained, it is not freed until the adapter is unconfigured. For a client adapter that has the rootvg device, the additional memory is not freed until the partition is rebooted.

Other software

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

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

AIXlink/X.25

AIXlink/X.25, Version 2.1

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

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

AIXlink/X.25, Version 1.1.5

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

AIX Fast Connect, Version 3.2

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

Quick start

To install AIX Fast Connect:

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

Communications Server for AIX, Version 6.1

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

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

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

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

Distributed Computing Environment (DCE)

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

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

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

Enterprise Identity Mapping (EIM)

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

Chapter 5. Documentation

AIX Information Center

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

- AIX 5L Version 5.3 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://publib.boulder.ibm.com/infocenter/pseries/v5r3/index.jsp>

Appendix A. AIX 5L Version 5.3 changes

New and Enhanced Functionality in AIX 5L Version 5.3

AIX 5L Version 5.3 delivers support for the System p POWER6™ based offerings.

Base Operating System (BOS)

64-bit system identifier

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

64-bit kernel

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

32-bit kernel

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

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

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

Long user names, group names, and path names

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

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

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

Notes:

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

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

System support

For information about supported and unsupported items, see Appendix B, “AIX 5L Version 5.3 unsupported devices,” on page 61.

Base functionality

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

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

- **System and I/O Support.** The 64-bit kernel limits support to 64-bit POWER-based systems, while the 32-bit kernel supports both 32-bit and 64-bit POWER-based systems. In addition, the 64-bit kernel does not support all I/O that is supported by the 32-bit kernel.
- **Application Support.** The 64-bit kernel supports both 32-bit and 64-bit applications. Application source and binaries are portable between AIX 5L Version 5.3 64-bit and 32-bit kernel systems, in the absence of any application dependencies on internal kernel details or on kernel extensions that are not supported under the 64-bit kernel but are supported under the 32-bit kernel.
 - **Binary Compatibility.** Binary compatibility is provided for 32-bit applications running on earlier versions of AIX on POWER-based systems, except for applications linked statically or applications dependent on undocumented or unsupported interfaces. In addition, some system file formats have changed, and 32-bit applications processing these files might need to be recompiled.
 - **Application Scalability.** AIX 5L Version 5.3 provides a more scalable application binary interface (ABI) for 64-bit applications. To take advantage of the scalability improvements to 64-bit programs, all 64-bit applications and libraries must be recompiled on AIX 5L Version 5.3. In addition, existing 32-bit kernel extensions and device drivers used by 64-bit applications might have to be modified in order to support the new 64-bit ABI.
- **Kernel Extensions.** Kernel extensions for the 64-bit kernel run in 64-bit mode and have the scalability of the larger kernel address space. Some kernel services available in the 32-bit kernel are no longer provided by the 64-bit kernel, so existing 32-bit kernel extensions may have to be ported in order to be used with the 64-bit kernel.

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

- **Dual-mode Kernel Extensions.** AIX 5L Version 5.3 supports dual-mode kernel extensions, which can be loaded by a common configuration method, regardless of which kernel is being used. A dual-mode kernel extension is an archive file that contains both the 64-bit and 32-bit versions of the kernel extension as members.
- **Installation and Enablement.** The 32-bit and 64-bit kernels are provided as part of the AIX 5L Version 5.3 base media and are installed on all supported hardware systems. The default kernel enabled during installation is dependent on the hardware system being installed. On POWER5 systems, the 64-bit kernel is enabled during base system installation. On all other systems, the 32-bit kernel is enabled. However, you can override this default option at installation time through the system installation panels. You can switch between the 32-bit and 64-bit kernels without reinstalling the operating system.
 1. Modify the `/usr/lib/boot/unix` directory and the `/unix` directory to be a symbolic link to the binary for the desired kernel.
 2. Run the `bosboot` command to write a new system boot image.
 3. Reboot the system.

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

JFS2 maximum file system support

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

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

JFS2 file system freeze and thaw feature

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

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

- **Command:**

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

- **API:**

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

JFS2 sync tunables

JFS2 sync tunables provide ways to limit the number of pages that are written to a file during sync operations. When there is heavy random I/O activity to a large file during a sync operation, all reads and writes from user programs to that file are blocked. When there are a large number of “dirty pages” in the file, the time required to complete the writes to disk can be lengthy.

Note: These tunables should only be used if application response is noticeably delayed when a sync occurs. The optimum setting for these tunables is dependent on memory size and I/O bandwidth. If tuning is required, start with a neutral setting of 256 for both tunables.

The following are definitions for JFS2 sync tunables:

j2_syncPageCount

Limits the number of modified pages that are scheduled to be written by a sync operation in one pass for a file. When this tunable is set, the file system will write the specified number of pages without blocking I/O to the rest of the file. The sync call will iterate on the write operation until all modified pages have been written.

j2_syncPageLimit

Overrides **j2_syncPageCount** when a threshold is reached to ensure that the sync completes.

JFS2 sync tunables are manipulated with the **ioo** command, and they can also be managed with **smit** or **wsm**.

Hints for sync tunable settings:

j2_syncPageCount
Default: 0
Range: 0-65536
Type: Dynamic
Unit: 4KB pages

When you are running an application that uses file system caching and does large numbers of random writes, you might need to adjust this setting to avoid lengthy application delays during sync operations. The recommended values are in the range of 256 to 1024. The default value of zero gives the normal sync behavior of writing all “dirty pages” in a single call. Small values for this tunable result in longer sync times and shorter delays in application response time. Larger values have the opposite effects.

j2_syncPageLimit
Default: 16
Range: 1-65536
Type: Dynamic
Unit: Numeric

Set this tunable when **j2_syncPageCount** is set, and increase it if the effect of the **j2_syncPageCount** change does not reduce application response times. The recommended values are in the range of 1 to 8000. The **j2_syncPageLimit** setting has no effect if the **j2_syncPageCount** setting is 0.

New memory allocation algorithm **MALLOCTYPE=watson**

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

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

The **-l** option removed from the **make** command

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

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

System performance recordings and reports

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

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

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

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

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

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

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

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

Reliability, availability, serviceability utilities

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

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

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

System dump

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

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

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

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

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

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

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

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

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

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

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

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

Advanced First Failure Data Capture features

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

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

```
smit ffdc
```

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

System trace

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

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

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

Lightweight Memory Trace

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

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

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

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

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

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

Component Trace

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

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

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

Run-Time Error Checking

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

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

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

RTEC features include:

1. Xmalloc debug

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

2. Excessive Interrupt Disablement Detection

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

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

Report these error logs to IBM service. The detailed report contains additional information including a stack traceback and LMT (trace) data that can be used by IBM to identify the source of the problem. Only one period of interrupt disablement that exceeds .5 seconds is logged per boot (default). Note that each of these error log entries might identify a unique potential problem. These error reports are persistently disabled if RTEC is globally disabled. On a per-boot basis, disablement detection can be disabled by running the following command:

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

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

```
disablement_checking_suspend  
disablement_checking_resume
```

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

Other RAS enhancements

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

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

AIX Web browser Mozilla Firefox 1.5 for AIX

AIX 5L Version 5.3 supports the Mozilla Firefox Web Browser Version 1.5.0.10 (or later) as the default Web browser for AIX. A version of the browser is available to be ordered on CD along with AIX. The latest version can be downloaded at no charge from the following Web site:

<http://www.ibm.com/servers/aix/browsers>

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

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

Installing Mozilla Firefox for AIX

Starting with AIX 5L Version 5.3 with the 5300-06 Technology Level, Mozilla Firefox for AIX can be installed as an option during the AIX Base Operating System installation process, or it can be installed later. All listed installation methods use the Mozilla Firefox installation bundle, which includes Mozilla Firefox and the required GNOME libraries.

The Mozilla Firefox installation process fails if the required GNOME libraries are not found. The required rpm filesets are listed. Remove the required rpm filesets if they were already installed for Mozilla.

Use one of the following installation methods:

- Install Mozilla Firefox using the following AIX Base Operation System (BOS) installation process:
 1. Select Firefox for installation during the AIX BOS installation process by selecting these options in the following order:
 - a. 2 = Change/Show Installation Settings and Install
 - b. 3 = More Options
 - c. 7 = Install More Software
 - d. 1 = Firefox (Firefox CD)

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

Configuring Mozilla Firefox as the browser for AIX documentation services

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

- Configure Mozilla Firefox using the following Configuration Assistant process:
 1. Start **configassist**.
 2. Select the Configure documentation server task.
 3. If Firefox is detected as already installed, select **Yes, use Firefox as the default browser**, and click **Next**.
- Configure Mozilla Firefox using the following **smit** or **smitty** process:
 1. Run the **smit change_default_browser** or the **smitty change_default_browser** command.
 2. Specify Firefox or the **/usr/bin/firefox** directory as the **defaultbrowser** launch command.

Migrating an existing Netscape Communicator Version 4 profile

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

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

License Use Management (LUM)

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

LUM Version 4

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

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

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

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

LUM Version 5

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

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

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

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

Exclusive resource sets

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

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

Multiple instances of AIX on a single root volume group

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

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

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

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

Multiple page size support

AIX 5L Version 5.3 with the 5300-05 Recommended Maintenance package includes 64-bit kernel support for virtual memory page sizes of 64 KB and 16 GB that are provided by POWER5+™ processors. These

virtual memory page sizes are supported in addition to the previously supported virtual memory page sizes of 4 KB and 16 MB. Using a larger virtual memory page size like 64 KB for an application's memory can significantly improve an application's performance and throughput due to hardware efficiencies associated with larger page sizes.

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

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

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

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

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

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

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

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

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

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

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

```
vmo -r -o vmm_mptest_support=0
```

USB flash drive support

Beginning with AIX 5L Version 5.3 with the 5300-09 Technology Level and AIX 6.1 with the 6100-02 Technology Level, USB flash drives are supported. To use these devices, install the following device package:

```
devices.usbif.08025002
```

The following USB flash drives are supported for use with AIX 5.3 with 5300-09 and AIX 6.1 with 6100-02:

- Lenovo USB2.0 Essential Memory Key, part number 41U4944
- Kingston 4GB DataTraveler, USB2.0
- SanDisk 4GB Cruzer Micro Flash Drive, USB2.0

Note: The AIX USB-host controller supports only USB1.1 on AIX 5.3 with 5300-09 and AIX 6.1 with 6100-02. All of the supported flash drives function with USB1.1.

These flash drives are configured with logical names such as flashdrive0 and flashdrive1, and they present both raw and block special files. For example, the raw special file for flashdrive0 is **/dev/rflashdrive0**, and the block special file is **/dev/flashdrive0**.

Only the ISO file system (read-only ISO9660) is supported on these drives. However, you can create a system backup on the drives by using the **tar** command, **cpio** command or the backup or restore archives. You can also use the **dd** command to add ISO images to the drives. These images can be used in read-only environments.

To make a flash drive available to AIX users, a root user must connect the drive to a system USB port and run the following command:

```
cfgmgr -l usb0
```

Note: Note: Use caution when you remove the flash drives from ports. If the drives are not properly closed or unmounted before you remove them, data on the drives can be corrupted.

After you remove the drives, they remain in the available state in the ODM until the root user runs the following command:

```
rmdev -l flashdrive<n>
```

When a drive is in the available state, you can reconnect it to the system, and it can be remounted or reopened. If a drive is disconnected from a system USB port while it is still open to a user, that drive will not be reusable until the user closes and reopens it. AIX 5.3 with 5300-09 and AIX 6.1 with 6100-02 do not include plug-and-play support for USB flash drives.

Communications, networking, and I/O

IP Security

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

Using Certificate Management System (CMS) with Java 1.4

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

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

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

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

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


The resulting stanza should show:

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

Write quit to save and close the file.

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

Asynchronous I/O fast path for CIO with JFS2

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

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

AIX Network Data Administration Facility

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

Internet Key-Exchange logging

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

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

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

RADIUS Server

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

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

Path MTU (PMTU) discovery

Path Maximum Transmission Unit (PMTU)-related information is now stored separately from the routing table, in a table called the PMTU table. Routes are no longer cloned for IPv4, and the **netstat -rn** command no longer displays PMTU values. A new **pmtu** command is provided to view the PMTU table. This command is used to display IPv4 and IPv6 entries and can also be used to delete a PMTU entry. A PMTU entry is added when the **route add** command runs with the **mtu** value specified. When a route is deleted, all PMTU entries using that route are also deleted. A **pmtu_expire** network option is provided to expire unused PMTU entries. The default value is 10 minutes.

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

AF_INET6 sockets

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

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

To preserve the former behavior (to continue receiving the IPv6 header followed by payload data), applications must now have their code modified to set the new **IPV6_AIXRAW_SOCKET** socket option on any **AF_INET6 IPPROTO_RAW** sockets and recompile.

Removal of support for devices

The following devices are not supported on AIX 5.3:

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

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

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

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

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

The devices.artic960 fileset

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

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

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

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

- **devices.artic960.rte**, IBM ARTIC960 Runtime Support

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

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

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

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

The **devices.pci.14108c00** fileset

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

Missing resource processing

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

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

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

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

IBM Tivoli Directory Server (LDAP)

To access the latest IBM Tivoli Directory Server 6.0 product information, go to the following Web site:
<http://www-306.ibm.com/software/tivoli/products/directory-server/req-aix.html>

Installation and configuration

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

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

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

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

- *Readme Addendum*

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

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

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

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

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

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

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

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

```
smitty aio
```

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

Administration

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

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

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

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

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

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

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

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

Internet Protocol (IP) over Fibre Channel

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

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

To set up the Fibre Channel networking capability between AIX and Thomson Grass Valley™ Media Servers (such as Profile PVS 1000). In addition, the following steps are recommended to establish the communication between the two host systems:

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

```
chdev -l fcnet0 -a enable_farp=no
```
2. If the Profile Media Server is set up to disable "checksum and re-transmit on error" (this is the default option), type the following command on the AIX command line, assuming `fc0` is the interface for this purpose:

```
ifconfig fc0 tcp_disable_cksum
```

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

```
ifconfig fc0 -tcp_disable_cksum
```

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

Sendmail, Version 8.13.4

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

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

Generic Routing Encapsulation

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

AIX iSCSI software initiator

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

For more information about configuring iSCSI, see the "iSCSI Software Initiator" section in the *AIX 5L Version 5.3 System Management Guide: Communications and Networks*. For more information about the **/etc/iscsi/targets** file, see the *AIX 5L Version 5.3 Files Reference*.

The iSCSI protocol driver implements RFC 3720, with the following limitations:

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

failure. The login redirection cannot occur between IPv4 and IPv6 networks. If the original login was on an IPv4 network address, the redirection must be to an IPv4 network address.

iSCSI Boot Support and IPv6 Support

AIX supports booting from an iSCSI attached disk when running on POWER6 or later processors and using the iSCSI software initiator and IPv4. An iSCSI boot does not support IPv6. An iSCSI boot is supported on Power Blade systems using either the iSCSI software initiator or the iSCSI TOE daughter card using IPv4.

The iSCSI TOE PCI adapter and the iSCSI TOE blade daughter card support only IPv4 for both boot and data disks.

When you are booting the iSCSI software initiator, ensure that the Ethernet network is configured so that the link comes up without delay. After the Ethernet link is enabled, the AIX iSCSI software initiator will attempt to contact the iSCSI target for approximately 30 seconds before declaring that the boot disk cannot be found and indicating "554 Unknown Boot Disk". Some Ethernet protocols, such as spanning tree protocols, might prevent the link from coming up in 30 seconds and will cause boot failures. Such protocols must be disabled or overridden on the Ethernet switch if they prevent the Ethernet link from coming up in less than 30 seconds.

Configurable IP Multipath Routing

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

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

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

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

Each of these policies works as follows:

Weighted Round-Robin (WRR)

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

Random (RND)

As the name suggests, a route is chosen at random from the multipath routing set.

Weighted Random (WRND)

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

Lowest Utilization (LUT)

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

Hash-Based (HSH)

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

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

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

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

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

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

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

Examples

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

```
==> netstat -rn
Destination      Gateway          Flags   Refs      Use  If   PMTU Exp Groups
Route tree for Protocol Family 2 (Internet):
default          9.3.149.65      UGc     0          0  en0   -  -  =>
default          10.10.10.3      UGc     0          0  en1   -  -
```

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

```

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

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

```

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

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

```

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

```

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

```

route set

```

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

```

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

```

The output is as follows:

```

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

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

```

Virtual SCSI client adapter

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

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

System management

AIX Network Data Administration Facility

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

Distributed Command Execution Manager (DCEM)

The installation of the CSM DCEM GUI (**csm.dcem.gui**) and the CSM DCEM Web-based System Manager application (**csm.dcem.websm**) packages is dependent upon the installation of the CSM Server. See “Cluster Systems Management” on page 25 for installation instructions.

dscrctl Command

Purpose

Allows the system administrator to read the current settings for the hardware streams mechanism and to set a system wide value for the Data Stream Control Register (DSCR). The Data Streams Control Register is privileged. It can be read or written only by the operating system.

For more information on the Data Stream Control facilities, refer to the POWER Instruction Set Architecture V2.05 publicly available from <http://www.power.org>.

Syntax

To query the characteristics of the hardware streams on the system:

```
dscrctl -q
```

To set the operating system default prefetch depth on the system temporarily (i.e, for the current session) or permanently (i.e after each reboot):

```
dscrctl [-n] [-b] -s <dscr_value>
```

To cancel a permanent setting of the operating system default prefetch depth at boot time:

```
dscrctl -c
```

Description

The **dscrctl -q** subcommand queries and displays the number of hardware streams, and the platform and operating system default prefetch depth. Any user can run this subcommand.

The **dscrctl -s** subcommand sets the operating system default prefetch depth. It requires root authority to run. This default can be changed either for the current session (-n), at boot time (-b) or both (-n -b).

The **dscrctl -c** option cancels the setting of the operating system default prefetch depth at boot time. This option removes the **dscrctl** command from the **/etc/inittab** file and thus takes effect after the next reboot.

Flags

- q** Query: Display the number of hardware streams supported by the platform and the values of the firmware and operating system default prefetch depth.
- c** Cancel: Cancels a permanent setting of the system default prefetch depth at boot time by removing the **dscrctl** command from the **/etc/inittab** file.
- n** Now: Used in conjunction with the -s flag, changes the runtime value of the operating system default prefetch depth. The change is not persistent across boot.
- b** Boot: Used in conjunction with the -s flag, makes change persistent across boots by adding a **dscrctl** command to the **/etc/inittab** file.
- s** *dscr_value*
Set: Defines the value for the new operating system default prefetch depth. The value is treated as a decimal number unless it starts with 0x in which case it is treated as hexadecimal.

Examples

1. Set the value of the operating system default prefetch depth to 13 for the current session:

```
# dscrctl -n -s 13
```

2. Display the current settings for the hardware streams mechanism:

```
# dscrctl -q
Current DSCR settings:
  number_of_streams = 16
  platform_default_pd = 0x5 (DPFD_DEEP)
  os_default_pd = 0xd (DSCR_SSE | DPFD_DEEP)
#
```

Related information

/etc/inittab

Enhanced nimadm command

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

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

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

Predefined XOPEN macros

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

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

It is set by default through `<standards.h>` if the strict standards scope feature macro is enabled. The **_ALL_SOURCE** macro enables all standards scopes and allows non-standard constructs to be visible as well.

- **_XOPEN_SOURCE**: when set to a value causes constructs defined in some UNIX specification to be visible

When **_XOPEN_SOURCE** is defined, each header defines or declares some identifiers, potentially conflicting with identifiers used by the application. The set of identifiers visible to the application consists of precisely those identifiers from the header pages of the included headers, as well as additional identifiers reserved for the implementation. There are various values set to correspond to various levels of UNIX specifications:

- **_XOPEN_SOURCE=1** is XPG4 or earlier
- **_XOPEN_SOURCE=500** is UNIX98
- **_XOPEN_SOURCE=600** is the correct feature macro to specify for strict UNIX03 support

- **_XOPEN_SOURCE_EXTENDED**: is UNIX95

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

Lightweight Infrastructure

The following AIX components rely on and incorporate Lightweight Infrastructure (LWI) v7.1:

- Electronic Service Agent™ (eSA)
- IBM Systems Director Console for AIX
- Common Agent Services (CAS) Agent

For information about licensing LWI v7.1, see the Program-unique Terms section of the AIX license information documentation.

Appendix B. AIX 5L Version 5.3 unsupported devices

Unsupported devices and machines

The following devices and machines are not supported:

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

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

Unsupported functions and filesets

The following functions and filesets are not supported:

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

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

Appendix C. Listing of filesets on the AIX media

AIX 5L for POWER Version 5.3 CD set

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

- **Volume 1:** Contains the minimal Base Operating System (BOS) software installed on every system, as well as all devices and both kernels, and the English message catalogs. To install a system with only "Volume 1" from the BOS menu, set Desktop to NONE, and change the default for Graphics bundle to No.
- **Volume 2:** Contains all of the software to install the Graphics bundle, and to set the Desktop to CDE. System management software, X11 software and Java software are on this CD.
- **Volume 3:** Contains additional software installed with the Graphics bundle, such as Director, Pegasus CIM software, OpenSSL (openssl), OpenSSH (openssh), Alternate Disk Install software, RSCT software, and the English man pages for libs and files.
- **Volume 4:** Contains pieces of Open_GL, PEX_PHIGS, LDAP, DB2 and additional ICU4C software.
- **Volume 5:** Contains software for installing AIX in other languages. Messages, help text, and locales for German, French, Italian and Japanese are included. Also contains AIX 3.2 Compatibility software.
- **Volume 6:** Contains software for installing AIX in other languages. Messages, help text, and locales for Catalan, Czech, Spanish, Hungarian, Polish, Brazilian Portuguese, Russian, and Slovakian are included.
- **Volume 7:** Contains software for installing AIX in other languages. Messages, help text, and locales for Korean, Simplified Chinese, and Traditional Chinese are included.
- **Volume 8:** Contains software to support bidirectional and complex text languages, including Arabic, Hebrew, Thai, Vietnamese, Hindi, Tamil, Telugu, Gujarati, Marathi, Kannada, and Malayalam. Chinese language support for Hong Kong and Singapore is also included. Also contains some additional performance software, printer software, RSCT LAPI messages, Cluster System Management software, and operating system updates.

Appendix D. Notices

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at www.ibm.com/legal/copytrade.shtml

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

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

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

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

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



Printed in U.S.A.

SC23-5201-08

