AIX Version 4.3 to 5L Migration Guide

- Latest enhancements in AIX 5L Version 5.1 and Version 5.2
- Planning and preparation hints and tips for your migration
- Step-by-step migration tasks by media and NIM

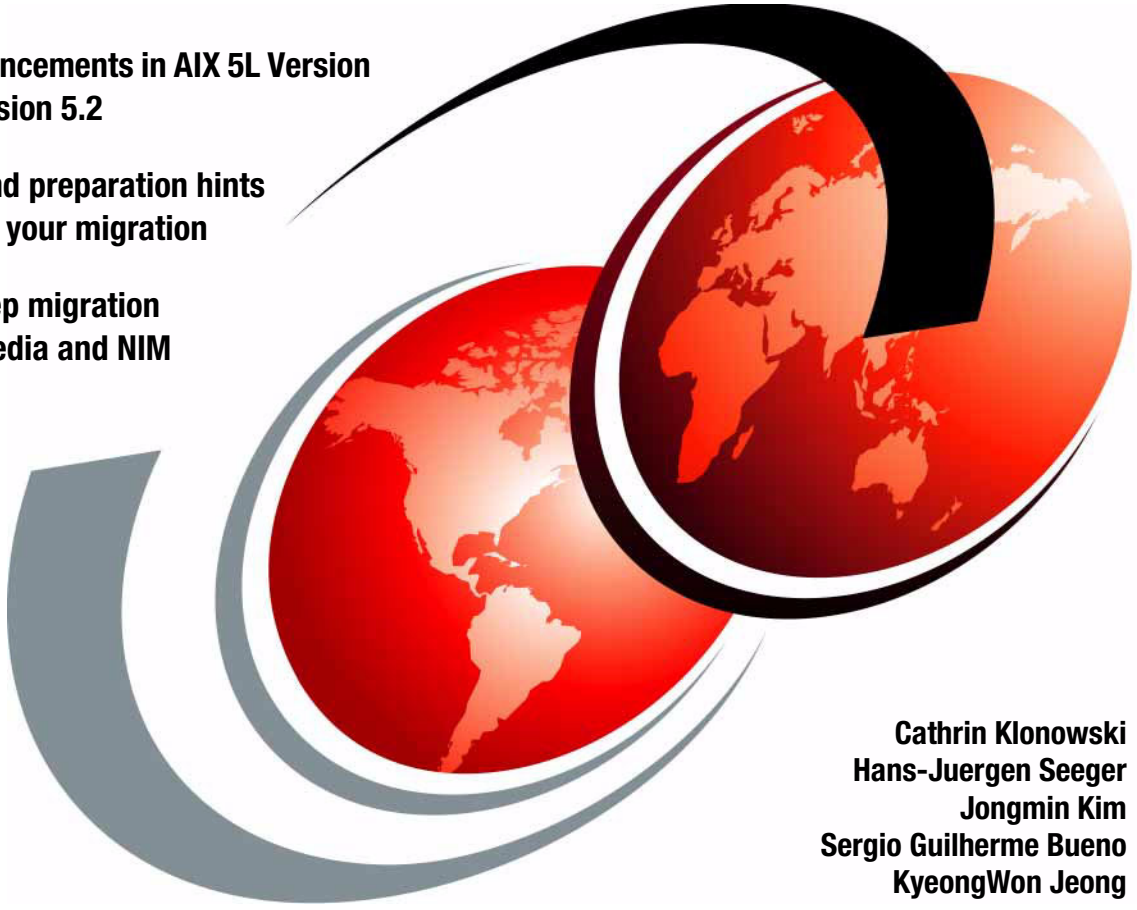

Cathrin Klonowski
Hans-Juergen Seeger
Jongmin Kim
Sergio Guilherme Bueno
KyeongWon Jeong

Redbooks

ibm.com/redbooks

International Technical Support Organization

AIX Version 4.3 to 5L Migration Guide

May 2003

SG24-6924-00

Note: Before using this information and the product it supports, read the information in “Notices” on page xi.

First Edition (May 2003)

This edition applies to IBM @server pSeries and RS/6000 Systems for use with the AIX 5L for POWER Version 5.1 Operating System (Program Number 5765-E61) or the AIX 5L for POWER Version 5.2 Operating System (Program Number 5765-E62), and is based on information available in February, 2003.

This document was updated on April 2, 2004.

Contents

Figures

Tables

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.

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.

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.

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.

COPYRIGHT LICENSE:
This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

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

AIX®
AIX 5L™
DB2®
DB2 Universal Database™


eServer™
ESCON®
GXT1000™
GXT150L™
GXT150M™
IBM®
Micro Channel®
PAL®
PowerPC®
pSeries™
PTX®
Redbooks™
Redbooks (logo)  ™
RS/6000®
SecureWay®
SOMobjects®
SP™
Tivoli®
VisualAge®
xSeries™
Xstation Manager®

The following terms are trademarks of other companies:

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation 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.

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

C-bus is a trademark of Corollary, Inc. in the United States, other countries, or both.

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

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

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

Preface

This IBM Redbook will provide AIX® system administrators with the key information they need to know in order to make a safe migration to AIX 5L™. It covers the steps that are required when planning, preparing, and actually migrating a system to the latest available level of AIX (currently AIX 5L Version 5.2) from AIX Version 4.3.

Before discussing the actual migration steps, we first point out the latest enhancements and new functions that have been delivered with AIX 5L Version 5.1 and Version 5.2 and why we recommend that you upgrade your system.

In the chapters that follow, we then explain how to plan, prepare, and perform the migration, focusing on providing hints, tips, and sources of additional information. We will explain detailed step-by-step actual migration by media and Network Installation Management (NIM).

This redbook also covers post-migration issues, including NIM master and alternate disk migration specific tasks. Lastly, we will summarize major applications issues when you migrate your system from AIX Version 4.3 to AIX 5L.

This redbook will help you design your solution to migrate your AIX system to the latest available version of AIX.

This redbook is essentially a follow-on to an earlier volume that described what was, at the time, the non-trivial task of upgrading to AIX Version 4.3 from the earlier Versions of AIX.

The team that wrote this redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization, Austin Center.

KyeongWon Jeong is a Consulting IT Specialist at the International Technical Support Organization, Austin Center. He writes extensively on AIX and education materials and teaches IBM® classes worldwide on all areas of AIX. Before joining the ITSO four years ago, he worked in IBM Global Learning Services in Korea as a Senior Education Specialist and was a class manager of all AIX classes for customers and interns. He has many years of teaching and

development experience. He is an IBM Certified Advanced Technical Expert - RS/6000® AIX.

Cathrin Klonowski is an IT Specialist from IBM Munich, Germany. She works in the Pre-Sales Technical Support Team in the Web Server Sales division of EMEA Central Region. Her areas of expertise and main responsibilities are pSeries™, AIX, and Workload Manager support to customers. She holds a Master's degree in Mathematics and Physics from York University, UK.

Hans-Juergen Seeger is a Software Engineer and works in IBM Global Services in the outsourcing section in Hannover, Germany. He has over nine years of experience with AIX and pSeries machines. He has worked at IBM for about 22 years. His areas of expertise include AIX system administration, AIX migration, IBM @server pSeries, and RS/6000 SP™ systems.

Jongmin Kim is a System Service Representative for IBM Korea in Seoul. He has worked in the ITS pSeries services team since 2000. He has three years of experience on AIX, RS/6000, and pSeries systems. He is an IBM @server Certified Specialist in pSeries, AIX System Administration, and p690 Technical Support. His main responsibility is customer support and the implementation of high-end RS/6000, pSeries, and SP solutions.

Sergio Guilherme Bueno is an AIX expert from IBM Global Services, Brazil. He has been working in AIX technical support since 1998. He is IBM AIX Certified in System Administration and System Support. He holds a degree in System Analyses from Faculdade de Administracao e Informatica (FAI) - Santa Rita do Sapucai, Brazil.

Thanks to the following people for their contributions to this project:

Julie Craft, Ann Wigginton, Edward Shvartsman, Susann Keohane, Stephen Nasypany, Mike Harrell
IBM Austin

Charlotte Brooks
IBM Almaden

Peter Fritz
IBM Austria

Glauco Jose Pinheiro and Leonardo Velloso Heitor
IBM Brazil

Armin Olaf Roell, Karsten Kutzer, Holger Holthoff, Stefan Kister, Helmut Droege, Georg Ember, Volker Haug
IBM Germany

Martin Bruegger
IBM Switzerland

Peter Nutt
IBM U.K.

Walter Orb and Rainer Staib
IBM SAP International Competence Center (ISICC)

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our Redbooks™ to be as helpful as possible. Send us your comments about this or other Redbooks in one of the following ways:

- Use the online **Contact us** review redbook form found at:

  ibm.com/redbooks

- Send your comments in an Internet note to:

  redbook@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. JN9B Building 003 Internal Zip 2834
  11400 Burnet Road
  Austin, Texas 78758-3493

1

Enhancements in AIX 5L Version 5.1

AIX 5L Version 5.1 represents the next generation of AIX. It is characterized by an increased level of integration, flexibility, and performance. This ensures support for the high demands of mission critical workloads. The new and enhanced features we will discuss in the following sections include:

- Resource Management and Control (RMC)
- Enhanced journaled filesystem (JFS2)
- Workload Manager monitoring tools
- Linux affinity

One of the main feature of AIX 5L Version 5.1 is logical partitioning (LPAR), which provides the capability to run multiple versions of the operating system on one physical piece of hardware. It is not in our interest to explain all advanced and enhanced new features of AIX 5L Version 5.1 in great detail in this chapter. We would like to mention the most interesting and most useful features, but refer you to the *AIX 5L Differences Guide Version 5.2 Edition*, SG24-5765 for a complete and detailed explanation.

1.1 Logical partitioning (LPAR)

The demand for systems that provide greater flexibility, in particular, the ability to divide the system into smaller partitions, each running its own copy of the operating system, keeps growing. Logical partitioning means splitting a hardware system into specific hardware boundaries (partitions) each running a separate image of the operating system, as shown in Figure 1-1.

Concept of logical partitions

The major advantage of a logical partition compared to a physical partition is the increased granularity and flexibility. This means that each logical partition consists of a set of CPUs, memory, and I/O slots from the general pool of resources. Several images of the operating system can be operated as independent business application servers. In other words, instead of having several small servers, each running a different application, you can have one powerful pSeries server with several logical partitions, each running a different application. The applications run just as they would on a single pSeries server.

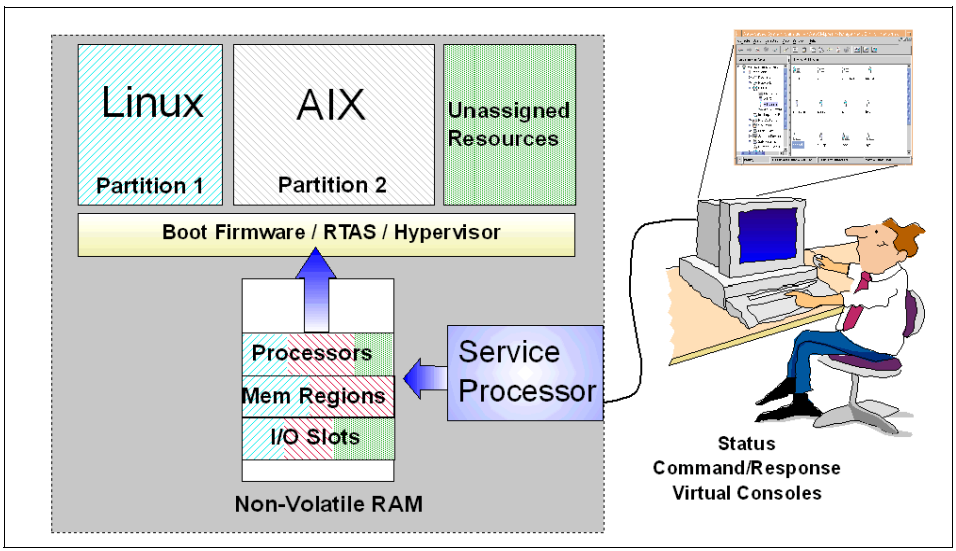

Figure 1-1 Logical partition concept

All POWER4 pSeries servers are capable of being logically partitioned. This includes IBM @server pSeries p630, p650, p655, p670, and p690. These servers are able to run either AIX 5L Version 5.1 or higher or Linux in one of these partitions completely independently of each other. The maximum number of partitions you can run concurrently depends on the pSeries server model and the physically installed resources on your server. These partition capable servers

are designed in a way to guarantee that there is no data access possible between partitions.

Benefits of using logical partitions

One important issue in logical partitioning is better resource utilization, particularly if you compare this concept to traditional UNIX servers. With the static concept of partitioning, as introduced in AIX 5L Version 5.1, you still need to reboot the system in order to redistribute your resources. This still gives you a lot of flexibility and ensures that you are not wasting resources when your workload changes temporarily or permanently.

Other advantages that come with logical partitioning are reduced floor space and reduced cost. This is due to the fact that you are making use of server consolidation aspects as well as simplified system management from a single point of control. This, however, means, as mentioned before, that you still have to manage each copy of the operating system independently. The access between partition data is permitted. Additionally, you can have various levels of the operating system on each partition, that is, each having different fixes and maintenance levels applied according to the application requirements. This means that problems on one partition do not affect any other partition with its applications. Only system wide problems will have an effect on all logical partitions.

For a more detailed discussion on the issue of logical partitioning, refer to *The Complete Partitioning Guide for IBM @server pSeries Servers,* SG24-7039. This publication explains all issues involved with static logical partitioning (as introduced with AIX 5L Version 5.1) as well as issues of dynamical logical partitioning, which is one of the most important features of AIX 5L Version 5.2.

1.2 System management

AIX 5L Version 5.1 offers many enhancements in system management. These include Resource Monitoring and Control (RMC), Web-based System Manager functionality, and system hang detection. Furthermore, there is the electronic license agreement, which is a new feature of AIX 5L Version 5.1. It is stored in the /usr/lib/objrepos/lag database file. This database is designed to integrate license information from non-IBM installation programs as well.

Resource Monitoring and Control

With AIX 5L Version 5.1 comes a new feature called Resource Monitoring and Control (RMC). It allows you to monitor the state of your system and to have responses and actions automatically taken. The system collects information about resource usage and performs predefined actions when a threshold value is

exceeded. It comes as part of the operating system at no additional charge and is installed and activated by default, but still needs to be configured.

From the technical point of view, RMC is a subset of the Reliable Scalable Cluster Technology (RSCT). This means that you can use RMC in a stand-alone or in a clustered environment as a single monitoring and management point. RMC gives you the ability to monitor system resources such as disk space, CPU usage, processor status, application process usage, and much more. AIX 5L Version 5.1 provides more than 80 predefined conditions that you can monitor. The concept is based on *conditions*, *responses*, and *actions*. The way it works is best explained using a simple example.

You first select a *condition* to be monitored. This could be the usage of your /usr file system, for example. You then define a threshold value, in our case, it is 90% file system usage. The *response* would be an informal notice, a warning, or an e-mail response. This means that the system administrator will get informed when the /usr file system usage exceeds 90%. Additionally, you can have the file system automatically expanded as the defined *action*. This is all done automatically by the operating system. To defined the three values, that is, condition, response, and action, you can use the Web-based System Manager, or command line.

RMC software

In Table 1-1 on page 5, we list the RSCT filesets that will be installed by default when installing or migrating to AIX 5L Version 5.1. Even though the RSCT daemons are active, you need to configure RMC on your system. That means you need to define conditions to be monitored as well as responses and actions taken by the system, as described in the example above. For additional information, refer to *A Practical Guide for Resource Monitoring and Control (RMC),* SG24-6615.

Table 1-1 RSCT filesets

Fileset name	Description
rsct.core.rmc	RSCT Resource Monitoring and Control
rsct.core.sr	RSCT Registry
rsct.core.errm	RSCT Event Response Resource Manager
rsct.core.auditrm	RSCT Audit Log Resource Manager
rsct.core.fsrm	RSCT File System Resource Manager
rsct.core.hostrm	RSCT Host Resource Manager
rsct.core.utils	RSCT Utilities
rsct.core.gui	RSCT Graphical User Interface
rsct.core.sec	RSCT Security

Web-based System Manager

The Web-based System Manager enables the system administrator to manage AIX machines locally or remotely. It is a graphical interface that provides many of the functionality you can obtain when using SMIT or command line. It is implemented using Java programming language, that is, Java 1.3 for AIX 5L Version 5.1.

There are several enhancements in the Web-based System Manager. These include dynamic user interfaces, Kerberos Version 5 integration in AIX, RMC interfaces and improvements in the management console, and SSL security options. Furthermore, the session log is introduced, which means all actions performed on any managed host will be logged, as well as successful and failed actions. Custom Tools, which include URL-based and shell executable-based applications, are integrated. Additionally, you can have customized environments for any user on any machine.

System hang detection

This new feature of AIX 5L Version 5.1 gives you the capability to recover your machine from certain system hangs. As with RMC, you configure in advance an appropriate action to be initiated in case of a predefined condition. This means that if your system hangs for one of the defined reasons, it can automatically be recovered. It can be configured using SMIT and relies on a daemon called shdaemon that has an entry in /etc/inittab. Several actions can be defined. They include generating error log entries, displaying warning messages, rebooting the system, and enabling a login shell with highest priority.

1.3 Networking

AIX 5L Version 5.1 provides new and enhanced features in the networking area as well. They ensure performance improvements and guaranteed system availability.

Virtual IP address (VIPA)

The virtual IP address support allows you to overcome application availability problems. With earlier versions of AIX, the application was bound to a real network interface. This was the only way to access the network. With VIPA, introduced with AIX 5L Version 5.1, the application is bound to a virtual IP address. This means that when the system detects a network failure, one of the other interfaces will be used for network traffic. The routing table will automatically be modified. This means improved availability for the application. This feature of virtual IP addressing is part of the bos.net.tcp.client fileset. The way it is implemented is that the virtual IP address is not associated with any particular network adapter. This new feature can be configured using the Web-based System Manager or SMIT.

Multipath Routing and Dead Gateway Detection (DGD)

Multipath Routing and Dead Gateway Detection are new features of the TCP/IP routing subsystem. They are implemented for IP Version 4 (IPV4) and IP Version 6 (IPV6) and provide an automatic selection of alternate network pathways. This means that your system can be configured with several routes to one destination. The multipath routing feature enables the system administrator to configure their systems for load balancing and failover.

The main advantage is that the alternate routes do not need to differ in their destination, netmask, or group ID list anymore. The mechanism to select one of the various routes is called round-robin, which is a cyclic mechanism. This means higher cost routes will only be used as backups. There are some limitations associated with the multipath routing mechanism. Refer to the *AIX 5L Differences Guide Version 5.2 Edition*, SG24-5765 for further information.

Dead Gateway Detection is used when one of the defined gateways is down. The system automatically recognizes this failure, adjusts the routing table, and re-routes the network traffic. This is something that was not possible with earlier releases of AIX. Up to now, you were not allowed or able to configure multiple routes to the same destination. With AIX 5L Version 5.1, there are two methods implemented. They include *Passive Dead Gateway Detection* and *Active Dead Gateway Detection*.

1.4 LVM and file system support

AIX 5L Version 5.1 introduces a variety of enhancements for the logical volume manager (LVM). These include new and enhanced commands like **redefinevg**, the **recreatevg** command, the /proc filesystem, and new LVM library support.

JFS versus JFS2

AIX 5L offers the flexibility of two kernels. This means that administrators can choose to run the system with either the 32-bit or the 64-bit kernel. Both kernel versions can run 32-bit and 64-bit applications. With AIX 5L Version 5.1, this 64-bit kernel enablement is linked with the creation of enhanced journaled file systems. This means that if your system is not 64-bit enabled, you cannot create JFS2 file systems. Furthermore, with a migration installation, the option to create JFS2 file systems is not available either. However, you can create these enhanced file systems once the migration process is finished manually.

JFS2 offers enhanced features that include extent based address structures, sorted directories, and dynamic space allocation for file system objects. The file system limit and size is increased up to 1 TB. This means that JFS2 can store much larger files, which gives the customer an improved, robust, and log-based file system alternative to JFS. For a comparison between JFS and JFS2, refer to Table 1-2.

Table 1-2 JFS versus JFS2

Function	JFS2	JFS
Fragment/Block Size	512-4096 Block Size	512-4096 fragments
Architectural Maximum File	1 PB[a]	64 GB
Architectural Maximum File System Size	4 PB	16 TB[b] (64-bit kernel) 1TB (32-bit kernel)
Maximum File Size Tested	1 TB	64 GB
Maximum File System size	1 TB	1 TB
Number of Inodes	Dynamic, limited by disk space	Fixed, set at file system creation
Directory Organization	B-tree	Linear
Online Defragmentation	Yes	Yes
Compression	No	Yes
Default Ownership at Creation	root.system	sys.sys
SGID of Default File Mode	SGID=off	SGID=on

Function	JFS2	JFS
Quotas	No	Yes
Extended ACL	Yes	Yes

a. PB stands for PetaBytes, which is equal to 1,048,576 GB.
b. TB stands for TeraBytes, which is equal to 1,024 GB.

Logical Volume Manager (LVM)

The LVM hot spare disk feature is introduced in AIX 5L Version 5.1 and allows you to define spare disks to a volume group that can be used in case of a disk failure. This is only successful in a mirrored environment that has no logically allocated partitions on the spare disk. The other requirement is that the spare disk must have at least the storage capacity of the smallest disk of that volume group.

One of the new commands that will improve your disk performance is the **migratelp** command. You use it when you have several logical partitions with heavy usage on one physical disk that you want to be balanced out over all available disks. With the **mkvg** command, the correct physical partition size is automatically determined when you create a new volume group. When you create or enlarge a logical volume using the **mklv** or **extendlv** commands, you can now specify the size in KB, MB, and GB rather than specifying the number of partitions.

When you migrate your system to AIX 5L Version 5.1, you will notice that there is a /proc file system automatically created on your system. It is a pseudo file system that maps processes and kernel structures to corresponding files. For further information on this file system and on other LVM enhancements, refer to the *AIX 5L Differences Guide Version 5.2 Edition*, SG24-5765.

1.5 Security enhancements

AIX 5L Version 5.1 provides various new and enhanced security features. These include enhanced Java security and Lightweight Directory Access Protocol (LDAP) enhancements. Furthermore, Kerberos Version 5 support can be used as an authentication loadable module for user logins. It is used for both the authentication and secure communication aspects of the client and server application.

Java security enhancements

The Java security enhancements are implemented through several Application Programming Interfaces (APIs), which give you the option to develop additional

Java applications. They are provided with Public-Key Cryptography Standards (PKCS) or Java Secure Sockets Extensions (JSSE). The latter feature is used to ensure secure Internet communication. In Table 1-3, we list enhanced Java filesets as 32-bit and 64-bit versions.

Table 1-3 Enhanced Java filesets

Java security enhancements	32-bit filesets	64-bit filesets
Certificate Management Protocol	Java130.cmp-us	Java130_64.cmp-us
Java Cryptography Extension	Java130.jce-us	Java130_64.jce-us
Java Secure Sockets Extension	Java130.jsse-us	Java130_64.jsse-us
Public-Key Cryptography Standards	Java130.pkcs-us	Java130_64.pkcs-us

LDAP name resolution

The Lightweight Directory Access Protocol (LDAP) provides a method for accessing and updating information in a directory. The name resolution mechanism has been enhanced with AIX 5L Version 5.1. It is extended to resolve host names through LDAP. In previous AIX releases, the name resolution mechanism only resolved host names through the Domain Name System (DNS), Network information Services (NIS), and the local /etc/hosts file.

An administrator can specify the ordering of name resolution services in the /etc/netsvc.conf and /etc/irs.conf files by setting the NSORDER environment variable to `local,ldap,bind,nis`. Furthermore, AIX 5L Version 5.1 provides an LDAP security plug-in, which audits the LDAP security information server.

1.6 Workload Manager (WLM)

AIX 5L Version 5.1 Workload Manager provides additional resource control and enhanced resource monitoring features. As explained in “Concept of logical partitions” on page 2, a new hardware architecture with POWER4 and the ability to use logical partitions on a powerful pSeries server lead to server consolidation. Due to even more powerful processors and enhanced bandwidth for some applications, the minimum requirements for a logical partition cannot be used to the full extent, which leads to resource waste. To ensure you are using the available system resources to their full extent, you can run several applications on the same system and on the same partition. That is where you can benefit from using AIX Workload Manager. You can control the resource usage in times of high server usage, that is, when there is high competition for resources between applications.

WLM provides isolation between user communities with very different system behavior. It is an operating feature that was introduced in AIX Version 4.3 and further enhanced with AIX 5L and can be used without additional charge. It is automatically installed on your system, but needs to be configured according to your environmental needs. One of the new features is the ability to manage not just CPU time and memory but also disk I/O bandwidth. Additionally, you can, independently of each item, regulate CPU, memory, and disk I/O, which means you can actively monitor CPU usage, for example, and not monitor memory and disk I/O at all or only in passive mode.

The configurations are dynamic, which means there is no need to stop and start WLM after each reconfiguration. You need to run an update command, though. The variety of ways that the resources can be managed has increased, as well as the possible number of classes and subclasses you can configure. Furthermore, graphical resource monitoring using the AIX Performance Toolbox have been introduced, as we describe in “WLM performance tools” on page 10.

Resource sets

A different way of using Workload Manager is using resource sets. This means that you can bind your application to a certain processor set. This only makes sense though if you have other applications that are not assigned to any set of processors. This means that if the major application that is processor bound does not use the resources, other applications can benefit from the available CPU resources. It is a different way to use WLM, because this way you assign your application to a certain resource, which means you restrict its use to only this resource even if others are free. It is a very useful concept in many customer environments. However, the way WLM works in general is that your application gets CPU time rather than having full CPUs assigned, which guarantees a finer granularity. Depending on your environment, you might use the general approach or resource sets or even a combination of both.

WLM performance tools

There are base AIX performance tools for WLM that are offered through basic LPPs. These include the **ps**, **wlmstat**, **svmon**, **topas**, and **wlmmon** commands. They are updated to monitor and display WLM-related data in AIX 5L. To benefit from these performance monitoring tools, you need to install the perfagent.tools fileset.

Additionally, there are Performance Toolbox enablements for WLM, which include **xmperf** and **wlmperf**. These are more advanced graphical interfaces for WLM monitoring. It also enables you to monitor the system and WLM activities graphically over a time period and to produce data reports. These can be used to analyze and save WLM data of CPU usage, memory, and disk I/O activity on your configured WLM classes over long periods of time (up to a year).

1.7 Linux affinity

AIX 5L Version 5.1 has a strong affinity with Linux. This ensures that you can compile and run your Linux applications on an AIX pSeries server. It is achieved by introducing the AIX Toolbox for Linux Applications and by integrating the Linux development environment into AIX libraries. This supports the customers that port their applications and use the features of both operating systems, AIX and Linux, at the same time. It also enables you to use basic Linux commands like **tar**, **gzip**, and **gunzip**, which will be installed in the /opt/freeware/bin directory.

AIX Toolbox for Linux Applications

The AIX Toolbox for Linux Applications contains a few hundred applications and software, which include system shells, application development tools like compilers, and Linux desktop environments like GNOME and KDE. These applications have been recompiled to be used on your AIX server. Table 1-4 shows a more detailed list of software, which is included in the AIX 5L Version 5.1 Toolbox for Linux Applications. The number and variety of these compiled applications keeps growing. For updated information, refer to the following URL:

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

Table 1-4 Software of AIX Toolbox for Linux Applications

Classification	Software
Application development	gcc, g++, gdb, rpm, cvs, automake, autoconf, libtool, bison, flex, gettext
Desktop environments	GNOME, KDE
GNU base utilities	gawk, m4, indent, sed, tar, diffutils, fileutils, findutils, textutils, grep, sh-utils
Programming language	guile, python, tcl/tk, rep-gtk
System utilities	emacs, vim, bzip2, gzip, git, elm, ncftp, rsync, wget, lsof, less, samba, zip, unzip, zoo
Graphics applications	ImageMagick, transfig, xfig, xpdf, ghostscript, gv,mpage, Gimp
Libraries	ncurses, readline, libtiff, libpng, libjpeg, slang, fnlib, db, gtk+, qt
System shells	bash2, tcsh, zsh
Window managers	enlightenment, sawfish

2

Enhancements in AIX 5L Version 5.2

IBM continues its strong commitment and investment in AIX. It is the strategic UNIX operating system for pSeries. AIX 5L Version 5.2 is robust and extremely reliable and has its main focus on the following areas:

- Performance and Scalability
- Continuous Availability
- Workload Management
- Infrastructure Management

One of the main features of AIX 5L Version 5.2 is Dynamic Logical Partition Support, improved throughput and resource utilization through increased Workload Management flexibility, and Cluster Systems Management Support. It offers support for large SMP server environments and high performance computing workloads through some scalability and storage enhancements. These include the following:

- Process memory affinity
- Multipath I/O for parallel SCSI disks
- JFS2 support for 16 TB files and file systems
- Enhanced LVM support for split mirror backups

AIX 5L Version 5.2 has implemented improved tools for managing your systems environment. Some Network Install Management (NIM) enhancements are introduced, as well as the possibility to do an alternate disk migration installation. Making a Linux Client for Web-based System Manager available underlines our strong commitment to Linux as the complementary operating system to AIX.

Application binary compatibility across all AIX 5L releases, the 32/64-bit functionality combined with advanced scalability, performance, RAS, and leading security features make it an attractive choice for customers to migrate their systems to AIX 5L Version 5.2. An overview of enhancements in AIX 5L Version 5.2 is given in the this chapter. For detailed information on these new and improved technology features, refer to the redbook *AIX 5L Differences Guide Version 5.2 Edition,* SG24-5765.

2.1 Self-optimizing functionality

AIX 5L Version 5.2 offers self-healing and self-optimizing functionality, such as dynamic CPU deallocation and reconfiguration. The system adapts automatically and dynamically to changes in the environment and reacts to disruptions appropriately.

DLPAR/DR

The concept of Dynamic Logical Partitioning (DLPAR) as a powerful new functionality was introduced in AIX 5L Version 5.2. It provides flexibility in dealing with changing workload demands and enables the system administrator to add, move, and remove system resources like processors, memory, and I/O slots from one partition to another without the need to reboot the system or any partitions.

There are many benefits to using DLPAR. You can move processors from a test to a production partition in periods of peak demands or you can move memory to a partition that is doing excessive paging. Another possibility of using the DLPAR functionality is to move an infrequently used I/O device such as a CD-ROM between partitions. The concept of DLPAR is shown and explained in Figure 2-1. It emphasizes that all Dynamic Reconfiguration (DR) requests will be initiated from the Hardware Management Console (HMC).

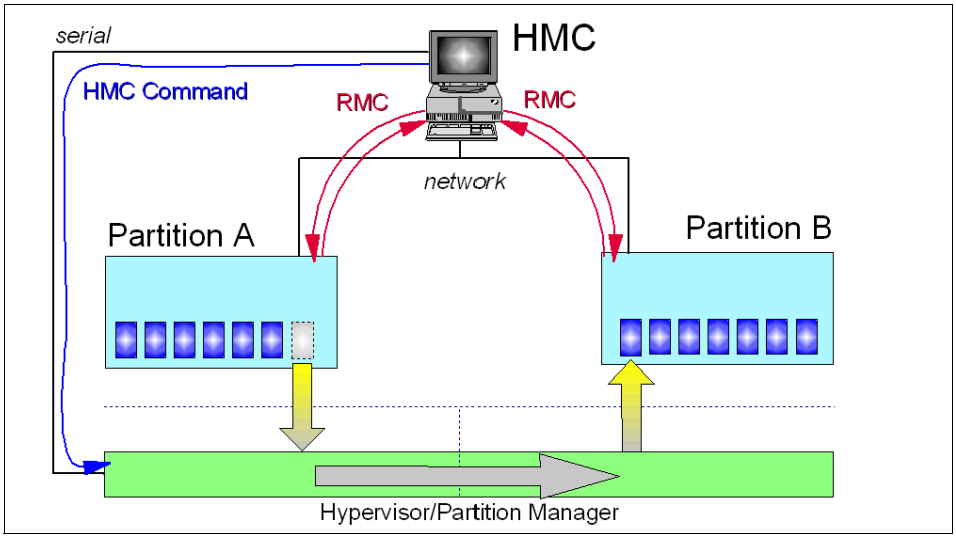

Figure 2-1 Concept of Dynamic Logical Partitioning (DLPAR)

In this example, we want to move a processor from partition A to partition B. We choose, on the HMC, the **Dynamic Logical Partitioning** menu for partition A.

We then check the box **Moving resource to a partition** and define partition B as the destination partition. This is all the administrator has to do. The rest is done by the operating system itself automatically and works as follows.

The HMC informs partition A about the request to remove one processor from partition A. The resource will be freed by scheduling all processes to other CPUs and handed over to the hypervisor, which acts as a partition manager. Partition A informs the HMC about the successful resource reduction. It has now six processors left. Through a serial connection, the hypervisor gets reconfigured from the HMC. This means that the free processor will be configured in the partition definition of partition B. This LPAR is then informed by the HMC about the request to add a processor. Partition B configures the processor and reports back to the HMC that the addition of the resource has finished successfully. Processes running in partition B can now be scheduled to use this additional processor in addition to the CPU resources that existed before. This ensures a performance improvement in times of peak workload demands.

Depending on your partition profile and what the required, desired, and maximum values are for the partition in terms of CPUs, it is possible to move more than one resource at a time. This works similarly for adding, moving, and removing memory. For I/O adapters, the dynamic reconfiguration is not fully integrated in the dynamic reconfiguration framework. Before you can move I/O adapters, you have to manually deconfigure the I/O slot by first using AIX commands on that partition. For a detailed description, including screen shots, refer to the redbook *AIX 5L Differences Guide Version 5.2 Edition*, SG24-5765.

Most applications are not aware of the amount of resources on the system. They are called DLPAR-safe and are not affected by DLPAR operations. As expected, there are DLPAR-aware applications as well. Examples are databases, because they scale with the configuration, performance monitors, because they report resource statistics, and license manager. DLPAR-aware applications need to adjust to the use of the system resources. They need to handle changes in the system configuration.

The dynamic reconfiguration (DR) framework introduced in AIX 5L Version 5.2 is designed to be non-destructive, which ensures that at no stage of the dynamic reconfiguration process a failure of exchanging resources can cause harm to the operating system image. With DLPAR technology, there are additional features enabled, which include Dynamic Capacity Upgrade on Demand (DCUoD) and Dynamic CPU sparing.

Capacity Upgrade on Demand (CUoD)

The Capacity Upgrade on Demand feature offers a non-disruptive method of activating additional processors based on resource needs. This is built directly into certain models of IBM @server pSeries, such as the pSeries 670 and 690.

To meet increasing workload demands, processors can be enabled dynamically. You can activate capacity in processor pairs only. It is a permanent activation and cannot be turned off. A keyed authentication based on CUoD Capacity Cards is necessary.

Static CUoD is available in SMP mode with AIX 5L Version 5.2. This requires a reboot of the system to change the number of licensed resources. Dynamic CUoD is available in LPAR mode, which does not require a reboot of the system or the affected partitions. DLPAR ensures that this operation is accomplished with no negative impact on the system. The main advantage of the CUoD feature is that customers can adjust easily to increasing workloads without requiring new servers.

Dynamic CPU Guard and CPU sparing

We differentiate between CPU Guard and CPU sparing. Both operations are considered to be DR operations. A processor reaching a pre-determined threshold is reported by Open Firmware. CPU Guard dynamically removes this failing processor. CPU sparing replaces the faulty processor dynamically. The spare processors are CUoD CPUs that are not licensed or activated with CUoD activation mode.

In conclusion, if you want to use Dynamic CPU sparing, you need to have an IBM @server pSeries 690 or pSeries 670 in LPAR mode that has DLPAR-enabled firmware and is loaded with an appropriate CPU Capacity Card. The replacement of CPUs is transparent to the user and the applications because no notification of reconfiguration will be sent to CPU Guard-aware or DR-aware applications. Using CPU sparing, you ensure performance and system availability for your pSeries server.

2.2 Cluster Systems Management (CSM)

CSM for AIX 5L provides you with the ability to efficiently set up and manage a cluster of various servers running AIX or Linux. The advantage is that customers can manage these clustered servers from a central location, which is called the management server. This saves you from having to manage each system individually.

CSM 1.3 is the first release of CSM for AIX. The version number 1.3 of CSM for AIX equals the version number of CSM for Linux that is required if Linux nodes are to be managed within the cluster. CSM for AIX is intended for new clusters, for example, for Cluster 1600 customers, in the commercial and High Performance Computing (HPC) area. It will be replacing Parallel System Support Program (PSSP) as IBM's main clustering software.

Requirements

The management server must be a pSeries machine running AIX 5L Version 5.2. The machines within the cluster are called managed nodes and are supported as a heterogeneous environment. The managed nodes have to be RS/6000 or pSeries running AIX 5L Version 5.1 Maintenance Level 03 or later and/or xSeries™ nodes running Linux Red Hat 7.2/7.3. The cluster management server can communicate with all nodes but all nodes, may not be able to communicate with each other.

Any pSeries or RS/6000 system running AIX 5L Version 5.1 ML 03 or Version 5.2 is supported under CSM for AIX but without any hardware control. In order to benefit from the hardware control feature, you have to use systems that are controlled by the Hardware Management Console (HMC), such as p630, p650, p655, p670, or p690. We will give a short overview of the main concept and functionality of CSM in this chapter. For detailed information on this topic, refer to the redbook *An Introduction to CSM 1.3 for AIX 5L*, SG24-6859.

Cluster functionality

From the management server, you can perform basic tasks on the nodes in your cluster. This includes monitoring cluster-wide hardware and software states as well as the installation of the operating system and CSM for cluster nodes using the Network Installation Management (NIM) functionality. In addition, you can remotely request node power status using the **rpower** command, do a reboot, and use power on and off functions. Cluster security is implemented using host-based authentication with public/private keys. This ensures that user and process authentication can be configured and maintained from the management server in a secure way. Other important features are the dynamic grouping of nodes and cluster diagnostics. For cluster-wide diagnostic purposes, logs from groups of nodes or single nodes can be collected and analyzed at the central location, the management server.

Comparing the clustering software features of CSM 1.3 for AIX and PSSP, we also find additional features with CSM. One big advantage is the ability to manage Linux on xSeries nodes as well as AIX on pSeries nodes. A tighter integration with AIX is achieved through Web-based System Manager and through the conjunction with Reliable Scalable Cluster Technology (RSCT). This means that nodes and applications can be monitored and automated responses are run when defined conditions on certain nodes are achieved.

Limitations

Due to the fact that it is the first release, there are a few limitations with CSM Version 1.3. The SP switch interconnect technology as well as the HPC software stack, including GPFS, Loadleveler and ESSL, are not supported at this time. The intention is, however, to support these features in later releases of CSM.

2.3 Scalability and performance

AIX 5L Version 5.2 supports large SMP environments and high performance computing workloads. Through enhanced performance and scalability features, the throughput and resource usage of pSeries hardware is maximized. Additionally, you are able to manage complex combinations of systems and software.

Multipath I/O for parallel SCSI disks

Multipath I/O (MPIO) is another new feature provided with AIX 5L Version 5.2. The benefits of using MPIO are an improved performance, easier administration, and more reliability and availability. The way it works is that one MPIO device or hdisk has multiple paths through different adapters of one machine or LPAR. Hence, you are able to uniquely detect, configure, and manage a device on multiple physical paths. These paths can be enabled or disabled without having to unconfigure the device. The initial release, however, supports SCSI disks only.

EtherChannel backup adapter

EtherChannel provides increased bandwidth, scalability, and redundancy. This is achieved by combining the bandwidth of multiple Ethernet adapters. The network traffic is distributed over all adapters in the channel rather than just over one. To implement this, several Ethernet adapters are aggregated to form one virtual adapter which the system views as only one adapter. To increase the channel's bandwidth, you simply have to add more adapters, but all the adapters share the same physical MAC address.

In AIX 5L Version 5.2, support for an EtherChannel backup adapter for EtherChannel installations was introduced. The backup adapter is used exclusively if all the adapters in the EtherChannel fail. In this case, the backup adapter takes over the IP and MAC address of the channel. In addition, you can protect your system against a switch failure by attaching the backup adapter to a different Ethernet switch than the EtherChannel adapters. This, however, requires all adapters in the EtherChannel to be connected to the same switch.

There are still limitations on which Ethernet adapters will work for EtherChannel and EtherChannel backup. Check the *AIX 5L Version 5.2 Release Notes*, GI10-0739 at the following Web site for updated information:

`http://publib.boulder.ibm.com/pseries/aixgen/relnotes/52relnotes.current.html`

Performance tools and tuning framework

There are several enhancements in the performance monitoring tools with AIX 5L Version 5.2. These include the Xprofiler, a GUI-based profiling tool for analyzing parallel and serial applications, the **pstat** command, which displays

many system tables, and the new implementation of the **tprof** command. The **pstat** command was missing in AIX 5L Version 5.1, but has been ported from AIX Version 4.3.3 to AIX Version 5.2 with the same functionality. Additionally, the **curt** command for taking AIX traces and the **splat** command for Performance Lock Analysis are introduced in this new version of AIX.

Performance management of AIX 5L Version 5.2 has seen some major changes in order to make kernel tuning parameters more flexible and centralized. Performance parameters like **vmtune**, **schedtune**, **no**, and **nfso**, which used to be lost after each reboot, are now stored in a new /etc/tunables directory. All the tuning commands have been enhanced to a consistent syntax and interface. The **vmtune** command is replaced by the new **vmo** and **ioo** commands. The **schedtune** command is replaced by **schedo**. All three new commands are part of the fileset bos.perf.tune. One of the enhancements with AIX 5L Version 5.2 is that these new performance tuning commands are also supported by the Web-based System Manager and SMIT.

Memory affinity for MCM based systems

Memory affinity offers improved performance by using the memory attached to the local Multiple Chip Modules (MCMs) rather than to the remote MCMs. This gives faster access and higher bandwidth, especially for applications that have processes or threads that initialize a memory array. It is available in SMP mode only and has no effect in LPAR configurations. It is enabled by setting the new MEMORY_AFFINITY environment variable to MCM on the AIX 5L Version 5.2 system. This way, the MCM local memory allocation is specified.

Note: The general discussion applies to POWER4-based systems that use a GX-BUS architecture.

2.4 Storage management

AIX 5L Version 5.2 offers support for emerging storage environments. The enhancements include improved disk handling by the logical volume manager (LVM) and JFS2 support for 16 TB files and file systems.

Logical Volume Manager (LVM)

In AIX 5L Version 5.2, the execution performance of the most used LVM commands has been improved. Especially extra large volume groups (BigVGs) with Volume Group Descriptor Areas (VGDAs), which are four times that of the disks, used to cause performance problems. The reason is that all the LVM shell scripts read and write the meta data several times during command execution. These shell scripts are rewritten in C and performance has improved by reducing

the number of times the command reads from and writes to the VGDA. Examples of commands which are now in C language are **mkvg** and **extendvg**. In this last case, the command will extend a volume group with all disks rather than with one at a time. The user, however, will only see an improved performance but no difference in the function of the commands.

Additionally, AIX 5L Version 5.2 provides support for adding Physical Partitions (PPs) to a redundant array of independent disks (RAID) that has grown in size. This is an important issue that caused problems and wasted space in earlier releases of AIX. With AIX 5L Version 5.2, however, a system that uses a RAID can grow the size of the physical volumes without disturbing the use and availability of the system. An informational message is automatically generated when growth of an activated volume group is detected.

Another interesting new feature of AIX 5L Version 5.2 is the snapshot support for mirrored volume groups, which is implemented using the new commands **splitvg** and **joinvg**. You can split a mirrored copy of a full mirrored volume group into a snapshot volume group and use it as a point-in-time backup of a volume group. At a later time you can reintegrate the split-off disk back into the volume group using the **joinvg** command. The stale partitions will then be resynchronized so that the user sees the same data in the rejoined volume group that was in the original volume group.

JFS2

With AIX 5L Version 5.2, the customer has the choice between implementing the journaled file system (JFS) or the enhanced journaled file system (JFS2). JFS is the recommended file system for 32-bit environments. The alternative is to have JFS2 running, which offers 64-bit functionality and provides the capability to store much larger files than the existing JFS. For the deployment of advanced applications and databases, the efficient storage for JFS2 on 32-bit machines is still 1 TB. On a 64-bit machine running the 64-bit kernel, however, the file system size has increased to 16 TB.

Support for JFS2 and for several 64-bit version commands have been added with AIX 5L Version 5.2. Examples of these are the **filemon** and **fileplace** commands. In addition, the JFS2 snapshot image has been introduced which is used to create backups of the file system at a specified time. This allows you to look at and to get information about files or directories at the time of the snapshot.

2.5 Security and directory

There are many new security items and authentication enhancements in AIX 5L Version 5.2. These include Controlled Access Protection Profile (CAPP) certification as an evaluated and assured quality implementation, AIX loadable authentication modules, LDAP enhancements, and new cryptographic functions.

CAPP and EAL4+ security installation

In order to install common criteria security code prior to AIX 5L Version 5.2, you had to order a special security CD that replaced the normal AIX product CDs. In AIX 5L Version 5.2, there are new security options available to choose at installation time. This allows you to install the CAPP and Evaluation Assurance Level 4+ (EAL4+) with the Base Operating System (BOS) CDs. However, this option is available for a new and complete overwrite of the operating system only, not for a migration or preservation installation. The installation can either be done using media or using the Network Installation Management (NIM) environment. In this case, the NIM master has to be a CAPP/EAL4+ system as well. At the time of writing this redbook, once CAPP is installed, your system can no longer be a NIM client. This will be fixed in the future.

The CAPP specifies certain functional requirements for the system that are similar to the old TCSEC C2 standard. A system with C2 security level enforces a discretionary access control policy that controls access to objects and protects information. It requires defined mechanisms for user authentication, user authorization, and process auditing, for example. This security standard is especially important for government agencies and for the banking and financing sector. For detailed information on CAPP and EAL4+ installation, refer to the *AIX 5L Version 5.2 Security Guide* which can be found at the pSeries Information Center home page at:

`http://publib16.boulder.ibm.com/pseries/en_US/infocenter/base/aix52.htm`

When you choose to enable the CAPP/EAL4+ technology, you are limited in your install options because there are prerequisites. These include having the Trusted Computing Base (TCB), the 64-bit kernel, and JFS2 enabled. English is the only accepted language choice. The only valid desktop choices are CDE or NONE. If you set any of these requirements to disabled, the CAPP and EAL4+ is automatically deselected.

There are many, very specific requirements for the physical and organizational environment for a CAPP/EAL4+ system. These include, for example, that physical access to the systems is restricted to authorized administrators only. When communicating over insecure lines, additional security actions need to be applied, like encrypting the information. One other restriction is that communication to another non-CAPP/EAL4+ system is permitted. Web-based

System Manager, certain Performance Tools, and Java are unsupported and forbidden as well. For the latest information on CAPP/EAL4+, refer to the *AIX 5L Version 5.2 Release Notes*, GI10-0739 at:

http://publib.boulder.ibm.com/pseries/aixgen/relnotes/52relnotes.current.html

Pluggable Authentication Module (PAM) support

in AIX 5L Version 5.1, the PAM libraries and include files were supplied with the first Update CD. However, they were not integrated into the AIX Security Services. This has changed with AIX 5L Version 5.2. Applications using the AIX security libraries, for example, can now call PAM modules that offer a flexible mechanism for authenticating users. Additionally, applications could call AIX security services using the PAM framework, which is implemented using an AIX Authentication Load Module, called PAM, and the pam_aix module. This is conceptionally similar to the Kerberos 5 or DCE. Shipped with the AIX 5L Version 5.2 CDs, however, is only the pam_aix module. If you want to use AIX security services to PAM module authentication, you either create your own PAM modules using the PAM framework or get PAM modules from a third party.

LDAP, public key infrastructure, and other enhancements

The IBM SecureWay® Directory Version 3.2 implements the Lightweight Directory Access Protocol (LDAP), which is offered with AIX for free and can be described as a centralized directory-based account management. AIX Version 4.3 and Version 5.1 use a proprietary schema to store the user and group security attributes. This feature has been enhanced with AIX 5L Version 5.2, where three schemas are supported: AIX, RFC2307, and RFC2307AIX. The AIX schema, for example, is included to support LDAP installations prior to AIX 5L Version 5.2.

AIX 5L Version 5.2 additionally supports a Certificate Authentication Service with Public Key Infrastructure (PKI). These PKI certificates as well as Network Information Services (NIS) are stored in the IBM LDAP Directory server. The subsystem information in AIX 5L Version 5.2 will be collected under a common subtree to simplify the administration in a directory enabled environment.

Another enhancement of AIX 5L Version 5.2 which we would like to mention is the support of a cryptographically secure pseudo-random number generator (PRNG), which is really important for any kind of cryptographic application. The AIX Internet Key Exchange (IKE) components now use the PRNG to generate random numbers. To make it easier to manage the IKE tunnels, some SMIT menus have been added or are improved. In addition, AIX Expansion Pack AIX 5L Version 5.2 CD includes an expanded cryptographic library.

2.6 Workload Manager (WLM)

There are several enhancements in AIX 5L Version 5.2 Workload Manager that provide even more control over resource regulation. The limit for user-defined superclasses was increased from 27 to 64. In addition, you can have 61 subclasses per superclass. This number was increased from 10.

Time-based configuration sets

One of the WLM enhancements of AIX 5L Version 5.2 are time-based configuration sets, which allow you to define a set of configurations to be active at a specific time. It is not necessary to define time ranges which cover all hours of the day. During these time range gaps, the default configuration will be active.

A configuration set includes a *.times* file, which lists all defined time ranges and their associate configurations. Regular configurations must be created first and can then be assigned to a configuration set. Configurations created prior to AIX 5L Version 5.2 are compatible to be used within configuration sets.

Attribute value grouping

Another useful enhancement is the attribute value grouping. Instead of listing out all the values of a specific assignment rule, an attribute value grouping can be specified. This can be a group of users, a group of applications, or other groups of attributes. It is configuration specific, but can also be copied to other configurations. By default, the grouping file is not defined.

Total limits

Additionally, in AIX 5L Version 5.2, there are two types of resources that you can regulate by setting limits for the total consumption. One way is to limit the total resource consumption of a process by specifying the maximum CPU time limit, total disk I/O, or the total connect time, which is the time a login session in a class can remain active. The assigned values are absolute limits and should only be applied to processes that should not consume resources excessively at any time. If the limit for any resource is reached, the process will be terminated.

Another way to control the resource usage is to limit the number of threads, processes or login sessions at a WLM class level. These limits are strictly enforced, which means when a class has reached its limit for one of the resources, any attempt to create another instance of the resource will fail. Enhanced commands for class total limits are **`wlmstat -T`** and the **`wlmcntrl -T`** command, which enable or disable the specified limits.

Event notification

AIX 5L Version 5.2 also provides an additional level of granularity and reports alerts not only on a system-wide basis but also at the WLM class level. For this task, the WLM Resource Manager (WLMRM) has been developed. It allows resource monitoring and control subsystem (RMC) clients to monitor resources at the WLM class level. It is possible, just as with normal RMC, to define new conditions to be monitored and to specify responses for notification of the administrator.

2.7 Linux affinity

AIX 5L Version 5.2 offers an enhanced affinity to Linux. This is accomplished through additional applications that you can install with the AIX Toolbox for Linux Applications. There are more than 380 applications available, which makes it possible to compile and run Linux applications on AIX. This way, developers can easily port their applications and run them directly on AIX, gaining all the features and benefits of AIX 5L Version 5.2. Additional information on this topic can be found in *Linux Applications on pSeries*, SG24-6033.

To download the AIX Toolbox for Linux Applications, visit the following Web site:

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

Web-based System Manager Client for Linux

Another new feature of AIX 5L Version 5.2 is the ability to run a Web-based System Manager Client for Linux. The Web-based System Manager client can be used to remotely connect to the Hardware Management Console (HMC) over either the network or the Web-based System Manager PC client. This PC client can run on a Windows operating system-based system or as a new feature of AIX 5L Version 5.2 on a Linux operating system-based system. This added Linux client is supported and tested on Red Hat 7.2 or Red Hat 7.3 Linux. A description of how to install the Web-based System Manager Client for Linux is described in detail in the redbook *AIX 5L Differences Guide Version 5.2 Edition*, SG24-5765.

2.8 IBM eServer™ pSeries Information Center

With AIX 5L Version 5.2, the IBM @server pSeries Information Center, a well documented Web site, is introduced. It can be found at the following URL:

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

You can find information on all AIX 5L Version 5.2 related topics as well as on earlier AIX releases. There is also a link to the complete IBM @server pSeries

library. Furthermore, you can search a database for error message numbers and LEDs, which helps you to identify the cause of your problem more easily. In Figure 2-2, we display the main Web site of the IBM @server pSeries Information Center.

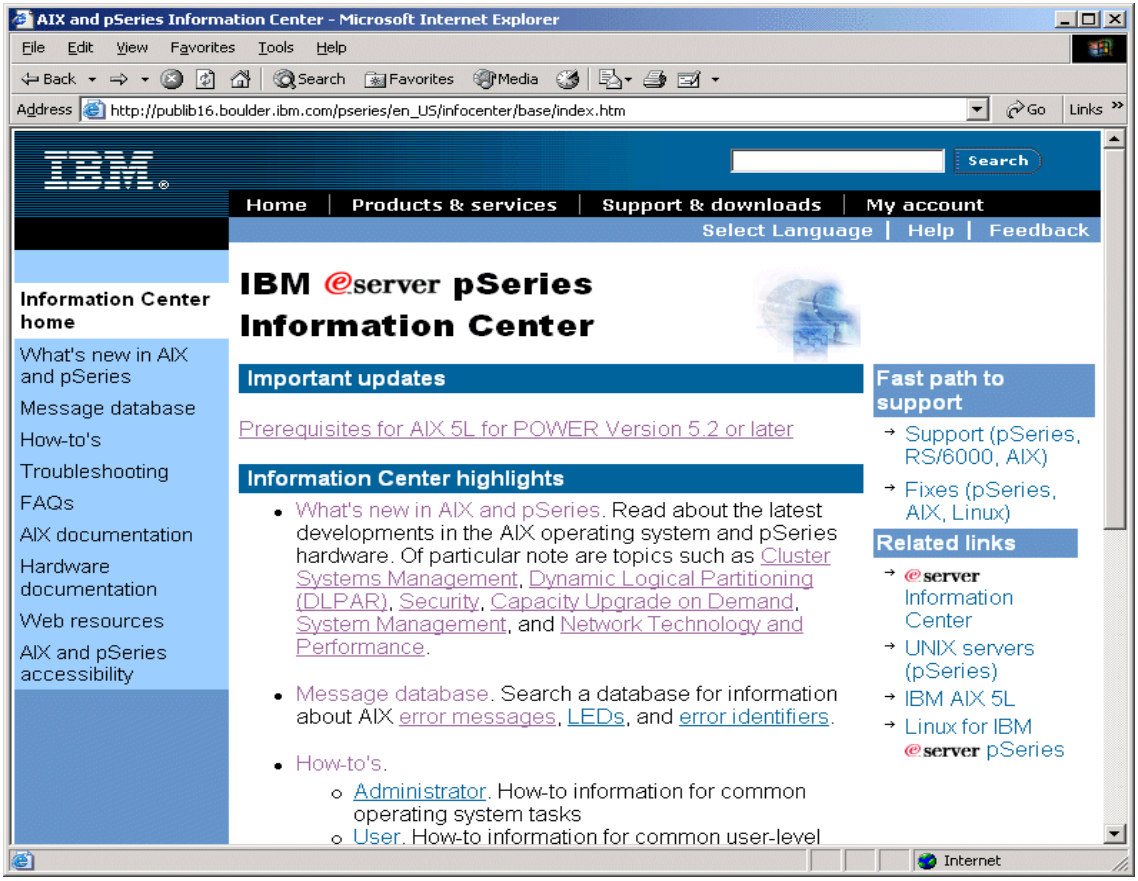

Figure 2-2 IBM eServer pSeries Information Center

To access this Web site from your AIX machine, you need to install the HTTP browser (set the DEFAULT_BROWSER environment variable). If you are using Netscape as the default browser, add the following line to the /etc/environment file:

```
DEFAULT_BROWSER=netscape
```

Then enter the **infocenter** command on the command line of your AIX server, which will display the main Web site of the Information center.

3

Seven reasons to migrate

After reading about the major enhancements in AIX 5L Version 5.1 and 5.2, you might ask yourself why should you read this chapter, after being given all the detailed technical information already? What else is there to come which could be a reason to migrate your AIX servers?

In this chapter, we will concentrate on a practical mapping of your application needs to the new capability of the operating system. We will not underline reasons why you should migrate your production environment to AIX 5L from a technical point of view but instead we want to view this migration issue from the customers point of view. This means that we include a discussion on what you as the customer expect from the operating system to fulfill your environmental needs and how this can be achieved with enhanced operating system functionality.

Even though we will not give a complete list of possible migration scenarios in this chapter, we will outline this mapping of imaginable environmental setups in 4.4, “Migration scenarios” on page 48. Your server environment might not match completely the described example scenarios, but we are certain you can complete the argument towards a migration to AIX 5L with your specific production environment in mind.

3.1 Software maintenance

The first important reason why a migration from AIX Version 4.3 to AIX 5L is necessary for all customers is the fact that support for AIX Version 4.3 will be withdrawn by the end of the year 2003. This means, after this deadline, you can continue to use AIX Version 4.3, but there will not be any fixes and maintenance levels available anymore. At the time of writing this redbook, there are three releases of AIX supported. These include AIX Version 4.3 and AIX 5L Version 5.1 and 5.2. The strategy is to guarantee to always have two fully maintained operating system versions available. From next year onwards, that is, 2004, the supported operating systems will be AIX 5L Version 5.1 and Version 5.2.

There are several reasons for this strategy. First of all, it ensures a permanent development of the software which you can use to have a more powerful and more flexible computing environment. In addition, it guarantees that we are able to deal and compete with the needs and requirements of customers environments.

Furthermore, major changes come with the development of a new hardware architecture. However, these POWER4 machines require AIX 5L Version 5.1 or later as the operating system. They are designed to only run on AIX 5L, which means you are not able to install and run AIX Version 4.3 on a POWER4 machine anymore. In order to benefit from the hardware enhancements and to participate in the new IT developments, you need to prepare your production systems. This gives another strong reason why it is time to migrate your server environment to AIX 5L.

Another important issue is related to security guidelines. With the end of maintenance AIX Version 4.3, there will be no more Program Temporary Fixes (PTFs) delivered for this AIX release. This includes security related fixes as well. If you have tight requirements on security aspects, which is essential for almost every customer environment, you have only one option to meet the requirements, that is, to migrate all machines in your production environment to AIX 5L and to be able to use the latest improvements and developments.

3.2 Networking

The modern computing scenarios are based on a complex and well functioning network structure. Years ago, computing systems took their power from the ability to do number crunching by using I/O devices for in- and output and terminals for control of operations only. Nowadays, the need for the user to access data at all times from anywhere in the world makes the network

availability one of the most important parts of the overall availability of the applications.

A strong advantage of the Wide Area Networks (WANs) is the possibility to connect two points via different paths. This means that when one route does not work because of a failure of one of the components in this route, network traffic along another path still provides the same service. Therefore, there are various Wide Area Network (WAN) technologies in place to ensure that this redundancy is available.

However, the way it works with AIX Version 4.3 is that the AIX machine is only aware of one way into the WAN. This ensures on one hand that a duplicate routing with all its disadvantages is avoided. This on the other hand means that it does not use the WAN technology to its full extent. When one route fails, even if there is another way into the WAN, the server is not able to find the alternative route automatically. It lacks the ability to analyze this network functionality by itself and to modify the routing table.

Among many other enhancements, AIX 5L Version 5.2 offers the possibility to have more than one default gateway defined. It can now automatically detect when a gateway is unavailable (Dead Gateway Detection). This means that the AIX server is able to select between a predefined set of alternatives for its routing. AIX 5L Version 5.2 is also capable of realizing when the missing gateway becomes available again. It can then switch the routing back to the originally defined value. It offers you the ability to define connections by different network providers and to switch between them automatically in case of failure.

The way it works with AIX Version 4.3 is that if one network provider asks for redundant network paths, in many cases different methods are selected. Usually the customer uses the same media or the same relay stations to carry the data. This means that when there is a problem with one provider, there is no connection possible, even if you have several backup lines from that provider. With the AIX 5L functionality, the system can switch automatically to the alternate provider. When both providers are available, you can even select the less expensive one using the cost based routing algorithm of AIX 5L.

In addition, applications keep growing and they keep getting more complex. This is even more the case when they use large and complex methods to structure themselves. Hence, they are running as a so called client-server application, which ensures that they are structured even during run time. The main characteristic is that the connection between the different parts runs over the local or distributed network interfaces. This means that if one of these interfaces fails, the application fails. This is the case even when the default route that the application is assigned to does not work because the associated interfaces have broken down, for example.

The solution to this problem is the assignment of virtual IP addresses, which were made available with AIX 5L. This is fully transparent to the application. In case of a network adapter failure, AIX 5L manages the rerouting and reassignment between virtual IP addresses.

3.3 Memory page size

Another important topic is memory. There is a rapidly growing number of applications that use an increasing amount of memory. This becomes necessary and possible with the increased functionality and the growing power of the processors. It is great for the customers, because it enables them to run complex applications which you could not run years ago. On the other hand, the computed results must be available in time frames that keep getting shorter. In many cases, being able to do all this online is another desired requirement. This shows how the requirements on the servers and the software are extremely high.

These advanced computing needs, however, require general changes in the design of memory layout and functionality. One restriction in accessing memory is the size of the pages the memory is divided into. It is also called the administrative unit for memory access and is generally limited to 4 KB. When a large amount of memory is handled, the system has to work with an administrative overhead of every 4 KB page.

With AIX 5L comes major improvements that can deal with these issues. You can define memory pools of 16 MB of size to allow selected applications to handle large memory resident data with an improved speed. The availability of large memory pools is not ensured for all applications, but for selected applications that really have the need to access the 16 MB memory pool benefit from this enhancement.

3.4 Disk storage

AIX or, in general, the UNIX operating system, provides a standard for organizing data on your disks. This leads to standardization on one hand but also to restrictions. As an example, years ago, the ability to create file systems, to be able to mount and unmount them and to have them under one common root directory, gave the operating system great flexibility. However, with increasing computing power, the demand for larger amounts of data, which needs to be handled and stored, is rapidly growing. To fulfill customers needs, the UNIX standard had to experience more enhancements.

This means that we are now able to have logical volumes spanning over the physically limited size of a hard disk. This is an important feature, but it creates

the need to have a faster recovery mechanism than the file system check after a system crash can provide as well. It can take hours, during which your system is unavailable, because all files of such a big system is checked after a crash. Hence, at this stage, journaled file systems were introduced to solve these issues.

The requirements to have faster access to the data stored on your disks and to make it easier and faster to back up your data keep growing in the same way as the requirements for more complex, more available, and more critical applications. To be able to follow these growing requirements for improved disk storage characteristics, AIX 5L, in particular Version 5.2, delivers a variety of enhancements.

When a file system spans over multiple disks, for example, a certain area of a particular disk is usually heavily used while different parts of the file system on other disks remain unused, AIX 5L Version 5.2 gives you the possibility to move dedicated partitions of that file system to other disks to spread the I/O activity to the less used disks.

The request for a large main storage area also requires the system to migrate this large amount of data into the storage area for processing. To support these increased needs in disk storage management, AIX 5L Version 5.2 provides the option to increase the transfer size of the data to the disk storage.

As far as the backup of application data residing on the disk storage is concerned, the size of data is growing while the window of availability for backup gets shorter and shorter. Even though the speed of the tape media has increased significantly in the last few years, it is not fast enough to directly make a backup onto the final backup media, so the application may have to be stopped.

Furthermore, the additional system load is not acceptable to many applications. Disk failures automatically have to be recovered, even during unattended operation. Because of these requirements, the disk storage management has been enhanced in AIX 5L Version 5.2. By using mirrored disks you have good availability until the moment when one of the mirrored disk fails. From that time on, you are operating without a mirror until the failing disk is exchanged. This may take several hours during which you do not have the high availability of mirrored disks. To cover this situation with AIX 5L Version 5.2, you can define one or more spare disks assigned to the mirrored set of disks. In case of a failure of a mirrored disk, the standby disk becomes part of the mirror and the time without a full mirror is therefore minimized.

From the system administration point of view, there are many enhancements included with AIX 5L that give the system administrator more flexibility to do system administration tasks without interfering with the running application. One

additional example is the ability to reduce the size of the swap space without rebooting the machine.

3.5 Hardware reliability and availability

With the growing number of users who work with UNIX applications, the demand for extended availability and 24/7 of service time is growing as well. This is especially important because UNIX servers are used more and more for business critical applications. This can be within one company or for Internet services. The use of computers for business gives opportunities to optimize processes to increase productivity and to gain money. On the other hand, a computer failure means a direct loss of money, even when the failure will be recovered by high availability solutions, because takeover takes time, during which the service is not available. This means that there are strong requirements for every customer to minimize the downtime or unavailability to the absolute minimum.

The current operating system architecture allows every processor to give any process only a little slice of its available computing time. This means that you have fine granularity, and in case of a CPU failure, only the work of this slice needs to be rescheduled. This is a hardware feature offered in combination with AIX 5L Version 5.2. This means that recovery points are available in a multiprocessor environment and the work can be done by another processor if one fails. This way, the processor failure would have no negative consequences for that process. Furthermore, this is achieved by integrating the processor taking over the whole processing environment. This is a complex task that needs to be supported by the hardware features as well as by the operating system. IBM @server pSeries POWER4 machines in connection with AIX 5L Version 5.2 provide these features.

Nevertheless, the new hardware flexibility supported with AIX 5L Version 5.2 is an advantage, and not only in the case of failure situations. The flexibility of many applications require different hardware configurations, depending on the workload produced or the state of the application. For example, applications that do monthly calculations require high computing power once a month for one or two days, whereas at other times of the month, the computing power could be available for other applications that are not time critical. Applications for planning purposes often use the weekend for intense calculations, but they use less computing power during the week days. The obsolete computing power during the week could be used for development or quality assurance systems.

The application software itself is changing faster, so static hardware configurations are not suitable anymore in many cases. There is a great need to easily change hardware configurations. This may require the reconfiguration of

processors, memory, or adapter cards. With AIX 5L Version 5.2 and the POWER4 hardware, these dynamic changes of hardware are possible while the application is running. There is no need to reboot the system or partition anymore.

3.6 Resource management with WLM

In the complex application scenarios of our customers, there are not just applications that always need more and more resources. A lot of smaller applications are required to collect the data that is needed by the major applications. These could be applications that are connected to terminal devices as batch- or card readers or that serve stock or production systems. In most cases, they are designed to run on small and isolated machines, but are needed for more than one business case. Because of this, they are installed multiple times on multiple small servers. Many of them are not designed to run multiple times on the same AIX machine.

Such applications mean cost problems when they have to be installed separately on the modern powerful AIX servers. One processor is the smallest unit you can assign to a partition at the time of the writing of this redbook. However, for such small applications, this is usually too powerful, as it is with memory. The minimum memory requirement is often much less than that of the modern operating system design. On the other hand, it is not always possible to install these applications multiple time on the same machine or partition because of the application design. This does not always allow you to run them more than once on the same operating system. This means that the more powerful the servers, the more problems you encounter with those applications.

In many customer environments, you use older and less powerful machines for this purpose. This is possible because the operating system releases up to AIX 5L Version 5.1 support all machines. The only thing that had to be upgraded was the amount of memory because of the operating system requirements. However, to benefit from the important enhancements of AIX 5L Version 5.2 you need to migrate to this operating system release. However, this means that you cannot use much of the old hardware anymore. The only supported machine type with AIX 5L Version 5.2 are CHRP machine types.

One solution for this problem is using the AIX Workload Manager (WLM), which manages system resources for multiple applications and ensures that they do not interfere with each other. This offers the possibility to run the same application multiple times with different base data. WLM ensures that the data of each application can only be accessed from this particular application. It means that each of these applications can run just as they would on a separate AIX

machine. It is possible even to install applications for multiple customers on the same machine.

With the Workload Manager functionality the amount of resource each application should get at a certain time can be defined. That means it is possible to independently define the following values to different applications or groups of applications:

- Amount of memory
- Amount of CPU time
- Amount of disk I/O

This ability of WLM gives you the chance to consolidate various different small applications onto one machine. One of the advantages is that you only maintain one operating system and hardware instead of multiple installations. It also protects your investigation into those types of applications, as they stay usable on the most recent computer architectures.

3.7 Security aspects

Security becomes more and more important in most customer environments. Security standards provided with the basic UNIX standard is not sufficient for the requirements of the actual applications, because the modern business processes are network driven. Each employee needs to access the company's databases to query product data and to be able to place orders directly. It becomes necessary to make the data accessible to a growing number of people and at the same time to protect the data from being modified from someone being unauthorized.

Furthermore, the customer's private data, such as credit card information, has to be protected from unauthorized access. The machines are placed in a public network access, but the data has to be protected and made as secure as possible. To assist with this task, AIX 5L has implemented several enhanced and new security functions.

There are more application programming interfaces (API) that enable dedicated applications to identify the users requesting access. More security methods are available to execute security tasks, like name resolution by using LDAP services or the usability of LDAP services in general. Another issue is the administration of your users, including user ID synchronization among a related group of servers, maintaining access permissions, and other similar tasks. When user administration exceeds a certain number of machines, a centralized administration is highly required. The use of LDAP services for authentication in AIX 5L allows you to do this task. It is not only possible to administrator the user

attributes in a centralized way, but also to restrict users to access certain machines.

For the network oriented operating systems such as AIX, there is not only the necessity to identify and authenticate users logging in, but also to perform tasks such as rsh or ftp, by doing remote access to several computers for the purpose of data exchange. Those tasks cannot follow the same login procedures as a user. The login has to take place unattended in the background. Mainly, the authentication leads to security problems as, in most cases, it is realized by files such as .rhosts doing IP address verification. This is not an acceptable way to authenticate any task requesting remote access. The use of keys that enable the authentication processes and users is an alternative method to solve this problem. AIX 5L Version 5.2 provides Public Key Infrastructure and Certificate Authentication Service as solutions to this requirement. The Pseudo-Random Number Generator (PRNG) is now supported in AIX 5L Version 5.2. You need random numbers for any kind of cryptographic applications, such as generating session keys or public key certificates. The AIX implementation of the PRNG is based on the Yarrow engine. For more information, refer to the following URL:

http://www.counterpane.com/yarrow.html

4

Planning and preparation

Before you start migrating your systems or your whole server environment, you have to take some time planning and preparing this step. This will ensure that you run into as few problems for your production environment as possible. For simple cases, the migration may just involve the straight-forward migration installation. For more complex cases, which involve third-party and in-house applications, we strongly recommend carefully planning the migration process.

In general, migrating your system and taking it out of production for this process is an important step for your server environment. It is a complex process, and if your company relies on these servers to run and function as desired, you need to carefully think about this step of migrating your environment from AIX Version 4.3 to AIX 5L and plan it in great detail. This way, a successful migration and a minimal downtime is guaranteed.

The migration planning process involves the following:

- Documenting your environment
- Mapping your environment
- Deciding on migration strategy
- Checking hardware and software requirements
- Backing up your system to prevent loss of data

We explain these issues and a few more in great detail in this chapter.

4.1 Hardware requirements

There are several hardware requirements for when you want to migrate your system from AIX Version 4.3 to AIX 5L. These are different for AIX 5L Version 5.1 and Version 5.2 and will be described in detail in 4.1.1, “Supported and unsupported hardware” on page 38. Additionally, there are required levels of firmware or microcode necessary for a server migration from AIX Version 4.3 to AIX 5L. We give information on these requirements in 4.1.2, “Microcode requirements” on page 39.

4.1.1 Supported and unsupported hardware

In this section, we give some general information on supported and unsupported hardware and device components. This differs between AIX 5L Version 5.1 and Version 5.2.

Hardware requirements for AIX 5L Version 5.1

AIX 5L Version 5.1 runs on most IBM @server pSeries and RS/6000 machines. These include, for example, pSeries 680 Model S85, p620, p640, p670, p690, and all models of p610, p630, and p660. The support for AIX 5L Version 5.1 includes RS/6000 Models 170, 260, and 270, Models H70, H80, and M80, as well as RS/6000 Models F80, S70, S80, and S7A. For a detailed list of unsupported devices and machines for AIX 5L Version 5.1, refer to Appendix A, “Hardware and fileset support with AIX 5L” on page 197 or read the latest *AIX 5L for POWER Version 5.1 Release Notes*, GI10-0729 at the following Web site:

`http://publib.boulder.ibm.com/pseries/aixgen/relnotes/current_relnotes.html`

There are several adapters that are not supported by AIX 5L Version 5.1 anymore. Some adapters are only supported in 32-bit kernel mode, but not in 64-bit kernel mode. These include graphics adapters like GX150P and GX120P as well as some ATM, Token-Ring, and Ethernet adapters.

Hardware requirements for AIX 5L Version 5.2

The IBM @server pSeries and RS/6000 machines that are listed above are supported with AIX 5L Version 5.2 as well. However, some machines have been withdrawn from support with this new version of AIX. This is the case for all Micro Channel® Architecture (MCA) and PowerPC® (PReP) architecture hardware. These include PReP Models 140 and 240, Scalable Parallel (SP) Nodes based on MCA architecture, and the PReP models F40, E20, and E30. This is, however, not a complete list of supported and unsupported hardware. Before planning your migration, check if your machine is a CHRP machine first by using the **`bootinfo -p`** command, because it is the only supported machine type. For confirmation, refer to Appendix A, “Hardware and fileset support with AIX 5L” on

page 197, which gives a complete list of machines for which support is withdrawn with AIX 5L Version 5.2. For other unsupported devices, such as graphics adapters, refer to the newest version of the *AIX 5L Version 5.2 Release Notes*, GI10-0739. You can find it through the IBM @server pSeries Information Center Web page, which has many more useful information on AIX and pSeries hardware related topics:

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

Unsupported and non-IBM hardware

If some of your hardware is in the list of unsupported devices or machines, you need to check whether or not it is only unsupported for the moment or if there are no plans to support it anymore. If that is the case, you need to either find a suitable replacement or manage your environment without this particular device.

The other important issue is that if you have non-IBM hardware or devices connected to your machines, you need to check with the manufacturer or supplier to see if the device is supported with AIX 5L to ensure your environment can function as normal after the migration to AIX 5L.

4.1.2 Microcode requirements

One important task to do on your system before starting a migration installation is an upgrade of the microcode level of your system and of all adapters and other devices. For this, visit the following Web site:

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

On this page, you can find a Microcode Update History on recent IBM @server pSeries and RS/6000 machines, including the latest microcode level. There is, in addition, a detailed download procedure for microcode updates. Before you download the microcode, check what level your machine is on and if an upgrade is required. You can get this information by typing the **lscfg -vp | grep alterable** command.

The output looks similar to Example 4-1.

Example 4-1 Determination of system microcode level

```
# lscfg -vp | grep alterable
ROM level.(alterable).......M2P01113
ROM level.(alterable).......CL010507
```

This works in the same way for adapters and other devices. An example of how to determine an adapter's microcode level is shown in Example 4-2. In this example, we want to determine the microcode level of an Ethernet adapter, ent0.

Example 4-2 Determination of adapter microcode level

```
# lscfg -l ent0 -v
DEVICE            LOCATION          DESCRIPTION
ent0              10-70             10/100 Mbps Ethernet PCI Adapter II
                                                      (1410ff01)
10/100 Mbps Ethernet PCI Adapter II:
Part Number.................09P5023
EC Level....................H10971A
Manufacture ID..............YL1021
Network Address.............0003474BDB7F
FRU Number..................09P5023
Product Specific.(PA).......A5204205
Product Specific.(OF).......SCU001
Device Specific.(YL)........P1-I4/E1
```

The `Product Specific.(OF)` line in Example 4-2 gives you the microcode level the ethernet adaper is currently at, that is, SCU001. In the next step, you need to compare this level with the currently available one and decide whether or not an microcode upgrade is needed.

Another important issue is the firmware level of your CD-ROM. If you do not have the most recent update, you may run into problems when you do a migration by media. Some AIX systems have problems booting from CD because the ROS Level and ID of the CD-ROM is lower than the required level. You can verify what level your machine's CD-ROM is using by running the **`lscfg -vl cd*`** command. If your ROS level and ID is less than or equal to 1_04 and your Part Number is 04N2964, a firmware upgrade of your CD-ROM is necessary.

When you upgrade your system's microcode, your system will be rebooted as part of the firmware installation process. When you plan your AIX migration, the firmware upgrade can either be done in a normal maintenance window or allow additional time in your migration schedule.

Even though this redbook describes the migration of AIX Version 4.3 systems to the AIX 5L operating system level, we like to make you aware that if you have a system running AIX 5L Version 5.1 with a Hardware Management Console (HMC) attached to it and you want to migrate the server to AIX 5L Version 5.2, you need to check the required HMC level in addition to the system and device microcode level. View the online *AIX 5L Version 5.2 Release Notes*, GI10-0739 for the latest information at the following Web site:

`http://publib.boulder.ibm.com/pseries/aixgen/relnotes/52relnotes.current.html`

Another issue is the software requirement for certain hardware. This information is important when you try to install or migrate your POWER4 system. In order to be able to boot from media, certain CD sets are necessary. For the IBM @server

pSeries server, which are not specified separately, the basic level of that AIX release works. The following examples are the requirements that are needed and necessary to know only if you move your applications from a POWER3 machine running AIX Version 4.3 to a POWER4 machine with AIX 5L:

- p630:
  - AIX 5L Version 5.1 5765-E61, dated 10/2002 (CD: LCD4-1061-04) or later.
  - For systems that are not used in LPAR mode, AIX 5L Version 5.1 at ML 02 plus APAR IY31315 is the minimum requirement in order to boot a non-LPAR system (AIX 5L Version 5.1 CD 5765-E61, dated 04/2002 (CD: LCD4-1061-03)).
  - AIX 5L Version 5.2: 5765-E62, initial CD-set (CD: LCD4-1133-00) or later.
  - AIX 5L Version 5.2 or later is required to support DLPAR.
- p650:
  - AIX 5L Version 5.1 5765-E61, dated 10/2002 (CD: LCD4-1061-04) and Update CD dated 12/2002 (CD: LCD4-1103-06) or later.
  - AIX 5L Version 5.2: 5765-E62, initial CD-set (CD: LCD4-1133-00) or later.
- p670:
  - AIX 5L Version 5.1 5765-E61, dated 10/2001 (CD: LCD4-1061-01) or later.
- p690:
  - AIX 5L Version 5.1 5765-E61, dated 10/2001 (CD: LCD4-1061-01) or later.

4.1.3 System requirements

There are some system requirements for AIX 5L Version 5.1 and Version 5.2. These include how much minimum memory and total disk storage is required. The minimum memory requirement for AIX 5L Version 5.1 is 64 MB; for Version 5.2, it is 128 MB. This means that if you want to migrate your system to AIX 5L you need to upgrade your AIX Version 4.3 system memory to at least 128 MB and then do the migration. This applies to a migration to AIX 5L Version 5.1 and Version 5.2. The reason for this task is that during the migration, there is a temporary need for more memory; if it is not available, you may have problems completing the migration successfully. Other approximate system requirements for AIX 5L are listed in Table 4-1 on page 42.

Table 4-1 System resource requirements

Resource type	AIX 5L Version 5.1	AIX 5L Version 5.2
Physical memory	64 MB	128 MB
Paging space	64 MB	512 MB
Disk space	664 MB	2.2 GB

If you need detailed information on how much allocated and used space is needed for /usr, /var, or any other directories for AIX Version 4.3 or AIX 5L Version 5.1 and Version 5.2, read the *AIX 5L for POWER Version 5.1 Release Notes*, GI10-0729 or *AIX 5L Version 5.2 Release Notes*, GI10-0739.

Another issue is that if your /tmp directory, for example, is smaller than the required size of 32 MB for creation of the AIX 5L boot image, it will automatically be increased during the migration installation. If you have an /opt directory on your AIX Version 4.3 system, you need to increase the size for the /opt directory to more than 3 MB before starting the migration process. If you do not complete this step, new logical volumes and file systems will be created during the migration for /opt or any file systems that had a mount point in this directory.

4.2 Software requirements

In addition to all the requirements mentioned in 4.1, “Hardware requirements” on page 38, there are some other issues you need to consider when planning your migration of your server environment to AIX 5L. Some may not be of any concern to you, but for completeness sake, they will all be mentioned in this section. If you require more detailed information on any of these issues, refer to the *AIX 5L for POWER Version 5.1 Release Notes*, GI10-0729 or *AIX 5L Version 5.2 Release Notes*, GI10-0739. Unless stated, these requirements or limitations apply to both versions of AIX 5L. For AIX 5L Version 5.1, go to:

`http://publib.boulder.ibm.com/pseries/aixgen/relnotes/51relnotes.current.html`

For AIX 5L Version 5.2, visit:

`http://publib.boulder.ibm.com/pseries/aixgen/relnotes/52relnotes.current.html`

If the level of xlC.rte is at 5.0.2.x or below on your AIX Version 4.3 system, it is necessary to apply APAR IY17981 before migration; otherwise, you might have problems with system boot after the migration process.

Some filesets will not be installed when you do a migration installation. These are known problems that can only be fixed manually at this time. If you are using the bos.clvm.enh fileset, you need to install it after the migration to AIX 5L Version

5.2 from media. If you are using the X11.loc.nl_BE.Dt.rte fileset, you need to perform a force install after migration. If you have the beta bos.pmapi fileset from the Performance Monitoring API (pmtoolkit) installed on your machine there are additional tasks to do before the migration to avoid boot failure of your machine. We do not describe the steps in this redbook. Refer to the *AIX 5L Version 5.2 Release Notes*, GI10-0739 for detailed information.

If you copy the content of the AIX CDs to a directory on the disk using the **bffcreate** command, be aware that the default target directory has changed with AIX 5L. It will now be copied to subdirectories of /usr/sys/inst.images, depending on whether you use the **installp** or **gencopy** command. This is important if you have scripts where the image location is hardcoded. In that case, you need to change them appropriately.

4.2.1 AIX Version 4.3 maintenance level requirements

When you plan your migration of AIX Version 4.3 to AIX 5L, one important issue is which maintenance level of AIX Version 4.3 is required as a basis or starting point. The latest maintenance level of AIX Version 4.3 available at this time is ML 10. You can upgrade your system prior to the migration to AIX 5L to this latest upgrade level (this is optional). There is no need for it because the migration installation code is designed in a way that it does not matter what release of AIX Version 4.3.3 you migrate from. In Chapter 5, “Migration by media” on page 79 and Chapter 6, “Migration by NIM” on page 105, we describe tested scenarios and will give you information on which maintenance level we use and which we have tested. In general, though, the migration of AIX Version 4.3 to AIX 5L should work without any particular maintenance level. The advantage is that no reboot and therefore no special maintenance window is necessary prior to the migration, as would be if an upgrade to the latest maintenance level was required.

4.2.2 Binary compatibility

There are various possible cases you need to look at to see whether or not your applications will run with or without recompilation after the migration to AIX 5L. The first case to consider is 32-bit applications running and compiled on AIX Version 4.3. These can be executed on AIX 5L without the need for recompilation.

Secondly, you cannot run a 64-bit application that was compiled under AIX Version 4.3 under AIX 5L without recompiling it. In addition, if you have a mixed server environment running AIX Version 4.3 and AIX 5L, you need to have two versions of this 64-bit application available. Furthermore, you need to select the correct binary for each platform, depending on whether you run AIX Version 4.3 or AIX 5L.

If your application was produced using AIX 5L on any of the 32-bit or 64-bit processor models, it will run on the 64-bit processor model without recompilation. The 32-bit application compiled under AIX 5L on either 32-bit or 64-bit processor models runs without recompilation on both models. For all other options, you need a recompilation.

The third case is the compatibility between AIX 5L releases. Across all AIX 5L releases, there is application binary compatibility and concurrent 32/64-bit functionality. This means that applications compiled and running on AIX 5L Version 5.1 can be executed on AIX 5L Version 5.2 without recompilation.

4.2.3 Availability of third-party and in-house applications

Verifying that third-party and in-house applications run on AIX 5L will be an important but time-consuming task. Make sure you plan enough time for application porting tasks. If you want to ensure that your system is functioning after the migration, you need to first make sure that your applications are binary compatible with AIX 5L prior to the migration process and that there are no licensing issues if the system is an application server. Depending on the complexity of your environment and the downtime requirements, we recommend you port your applications onto a test server and do some proper testing before you move your production environment to AIX 5L. Additionally, refer to your application documentation or provider to verify on which level of AIX your application is supported and licensed. The following Web site may useful for this issue:

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

Furthermore, you might want to check that all currently installed software is correctly entered in the Software Vital Product Database (SWVPD). To verify that all filesets have all required requisites and are completely installed, type the **`lppchk -v`** command.

4.2.4 Additional software products with AIX 5L

Many software products are specific to AIX Version 4.3 or AIX 5L release. Some may not work without deinstallation prior to the migration process. Examples include the AIX Toolbox for Linux Applications and the Performance Toolbox. Other software products, such as compiler and HACMP availability, is described in Chapter 8, “Applications” on page 179. An additional consideration is the topic of customized SMIT menus.

AIX Toolbox for Linux Applications

The AIX Toolbox for Linux Applications from your AIX Version 4.3 system may not run correctly after the migration to AIX 5L. You need to remove this software

from your system prior to the migration process if the level of the rpm.rte fileset is less than 3.0.5.20. You can remove this software by using the destroyrmps tool from the Toolbox CD. A sample of the commands to use is shown in Example 4-3.

Example 4-3 Deinstallation of AIX Toolbox for Linux Applications

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

The reason why a de- and reinstallation is necessary is that changes have been made to the shared library construction. You find additional information in the README.TXT file on the AIX Toolbox for Linux Applications CD. This applies to a migration of AIX Version 4.3 systems to both versions of AIX 5L.

Performance Toolbox and performance tools for AIX

When you migrate your system from AIX Version 4.3 to AIX 5L, certain files are removed from the system. These include /usr/lib/drivers, /usr/lib/microcode, /dev, and /usr/lib/methods. This means that software support for device drivers must be reinstalled. Additionally, some software products are removed from the system, which includes the Performance Tools functionality from Extended Commands.

Performance Toolbox (PTX®) Version 2 was withdrawn from marketing and service in December 2002. This means that at the moment you can work with PTX Version 3, which is supported on AIX Version 4.3 and AIX 5L. The new version, PTX 3.1, which will come out in 2003, will only be supported with AIX 5L.

You do not have to deinstall PTX prior to the migration process. PTX Version 3 can be installed over Version 2.2. Customers using AIX 5L Version 5.2 and PTX Version 3 should be aware that there is a coreq problem with one of the PTX perfmgr.analysis filesets, which is fixed by APAR IY38023.

Table 4-2 gives an overview of the supported versions of the Performance Toolbox on AIX Version 4 and Version 5.

Table 4-2 Support matrix for Performance Toolbox

		AIX operating system version		
		4.3.3	**5.1**	**5.2**
Performance Toolbox (PTX)	Version 2.2	Y	Y	Y
	Version 3	Y	Y	Y

Customized SMIT menus

If you have customized SMIT menus on your AIX Version 4.3 system, these will be functioning after the migration to AIX 5L as before. The migration installation keeps the SMIT ODM intact.

4.3 Migration strategy

There are several possibilities to upgrade your AIX Version 4.3 system to AIX 5L Version 5.2. All of them work perfectly fine and will be explained in detail in the following chapters. Which one will be the most attractive one for you really depends on your server environment. You might even find that no migration method suits all of your machines and that you have to use a combination of them. All migration methods certainly have advantages in some way, but some of them you might not be able to use because of company requirements, limitations on spare disks or machines, and downtime requirements. We are going to outline each method in detail, including advantages and drawbacks, in this section, and hope you can use the information to decide on your migration strategy.

4.3.1 New and complete overwrite

Doing a new and complete overwrite of your system means you are losing all existing data and information of your AIX Version 4.3 system. The target disk will be overwritten. Since it is not really a migration installation, it will not be explained in much detail in this redbook. However, in some customer environments, it may be a good choice to get a clean new AIX 5L system and maybe use it as Network Installation Management (NIM) master. One advantage is that you have a fresh installation with default configuration files only and that a overwrite installation is quicker than a migration. Another reason for doing a new and complete overwrite is that you might want to reassign your existing rootvg to fewer hard disks because it is spread out over several disks at the moment. In order to keep using the Trusted Computing Base (TCB) security features with your AIX 5L system, you need to do a overwrite installation as well.

There are, however, drawbacks to this installation method. The most important one as mentioned before is that you lose all existing data and configuration files. This means that you have to reconfigure your system, including TCP/IP and user information. This could be quite time consuming, depending on the size and complexity of your environment.

4.3.2 Migration installation

A migration installation is a method to upgrade your system from one release of AIX to the next. In our case, we are concerned about the migration of AIX

Version 4.3 systems to AIX 5L. When you insert an AIX 5L CD into the CD-ROM of the machine that currently runs AIX Version 4.3, you will automatically get, as the default installation option, the migration of the operating systems in the installation menu. This is one way of performing a migration installation. The other possibility is to use NIM. Both migration paths will be explained in detail in Chapter 5, “Migration by media” on page 79 and Chapter 6, “Migration by NIM” on page 105.

The advantage of a migration installation compared to a new and complete overwrite is that most filesets are preserved. This includes almost all directories, such as /home, /var, /usr, the root volume group, logical volumes, system configurations and previously installed software. The only file system that will be new after the migration is /tmp. You can easily avoid loosing information you have stored in this directory by copying the important information to another directory before the migration and move it back afterwards. Additionally, after the migration, you can import your user volume groups. It is probably the easiest way to upgrade your system to AIX 5L while maintaining all customized information and configuration. Another advantage, especially if you need to minimize the downtime of your system, is that there are fewer reconfigurations tasks to do after the migration.

One disadvantage that might be of concern in your server environment is the fact that you cannot create JFS2 file systems with a migration installation. This option is only possible for new and complete overwrite and preservation installations. Another limitation is that if you want to install graphics software support and other additional software that you can select with the installation menu, this option is not available at migration installation time. However, after the migration is completed, you can select additional filesets to install from various CDs or, alternatively, you can prepare software resources before migrating your system over the network and have them installed with the migration process already.

4.3.3 Alternate disk migration

One very attractive option to migrate your existing server environment from AIX Version 4.3 to AIX 5L with as little downtime as possible is the alternate disk migration installation. The only administrative task is to prepare the NIM server with the necessary filesets and then the administrator initiates the process of alternate disk migration by using one command only, the **nimadm** command. The procedure behind this command is a mixture of cloning your disk and using Network Installation Management (NIM). It is done by your server automatically and involves several steps. The starting point is to create a copy of the root volume group on a free disk in your system. Simultaneously, this disk gets migrated to AIX 5L while your original AIX Version 4.3 system is still running as normal.

The advantage is certainly that your down time is reduced because you migrate the system while it is still running and functioning. The other big advantage is that if the migration installation on your cloned rootvg fails, the changes made from the migration are only made to the copy of the rootvg. You can either clean up the failed migration or redo the whole process, including making a copy of the rootvg again. This is all done while your system is up and running without any major disturbance. You might have a small performance decrease on your active rootvg during the cloning due to increased disk I/O and CPU usage.

Despite all the advantages of this migration method, there are certain requirements and limitations. One important restriction has to do with the Trusted Computing Base (TCB). If your AIX Version 4.3 system has this option turned on, and you want the TCB enabled on your migrated system running AIX 5L as well, then you either need to do a new and complete overwrite or a 'normal' migration installation. Choosing to do the alternate disk migration installation is not possible, because TCB needs to access file metadata that is not visible over NFS. Furthermore, you need a machine set up as NIM master with AIX 5L in addition to your AIX Version 4.3 system.

An example alternate disk migration will be performed and explained in detail in 6.4, "NIM client migration" on page 137.

4.3.4 The mksysb installation

This method is called the mksysb install because you back up your system using the **mksysb** command and then restore it onto another system. This option is especially interesting for customers who have POWER3 systems and want to move their applications to a POWER4 machine, such as p630, p650, p655, p670, or p690. These POWER4 systems, however, require AIX 5L.

This means that one possibility is to do a fresh installation on your POWER4 partition and to transfer the user data only to this partition. Another possibility is to migrate your POWER3 system first from AIX Version 4.3 to AIX 5L using any of the described migration methods. Then you make a system backup and restore it onto a POWER4 machine. All you have left to do is the individual customizing after the restore of the mksysb install. A detailed discussion on this issue is described in Chapter 5, "Migration by media" on page 79. The same setup, but this time including any NIM related issues, is discussed in detail in 6.3, "NIM master migration" on page 117.

4.4 Migration scenarios

When preparing for migration to AIX 5L, there are a number of issues and questions that you need to consider. The first step is to get an overview of the

systems in your environment that need to be migrated. The next step is to make a decision on your migration strategy. To help you decide which migration method to use, we describe in 4.3, “Migration strategy” on page 46 different ways of migrating from AIX Version 4.3 to AIX 5L, including advantages and drawbacks of each method. All of these methods are described in detail, using our test example setups, in Chapter 5, “Migration by media” on page 79 and Chapter 6, “Migration by NIM” on page 105.

If you have several machines in your environment, we recommend using the Network Installation Management (NIM) method. In such an environment, it is the most time-saving method and ensures that customized scripts can be set up once and then used over and over again. Another advantage is that not every machine needs to have its own CD-ROM drive for the migration installation and it guarantees that all systems are installed with the same basic filesets. The convenience and advantages in productivity are considerable and easily justify the effort required to set up the NIM server.

A very attractive solution is a method called alternate disk migration. If you have tight requirements on the downtime of your server, it is a very good solution for your environment. The advantage is that you clone your rootvg and simultaneously migrate it from AIX Version 4.3 to AIX 5L while your server is up and running as normal. When your cloned and migrated root volume group is set up as your production environment requires, all it takes is a reboot to set the AIX 5L rootvg disk as the first boot device. This way, the downtime of your production server is reduced to the boot time of your AIX machine.

You may even find that none of these migration methods is ideal for your complex environment. If you have a heterogeneous environment with a mixture of servers and workstations from different vendors, then the upgrade strategy may be slightly more complicated. This means that you might end up using a combination of these methods, but as you can see in the following chapters, all the methods work if you carefully plan and prepare the migration process. This includes having supported hardware and software, all required filesets, a manageable time frame for the migration set up, a well-documented and backed up environment, and a recovery strategy or fallback options for the unlikely case that things go wrong with your AIX migration.

Another important issue is the time frame of the migration. You need to make the decision of when it is a good time to do the migration. This is necessary to ensure your operating system migration causes the least amount of disruption in your production environment. You may consider setting up a test migration environment first and do the actual migration after you have verified that the migration works without problems. You may also consider migrating all of the most important machines in your environment first, such as the NIM master or application server, and migrate all the other machines at a later time. If you

decide to migrate your environment using NIM, the NIM master must be migrated first, because it should run on the highest AIX level of all machines in the NIM environment.

And finally, you need to think about your applications. Are they supported by AIX 5L? Do they need to be upgraded for the new version of AIX? If that is the case, or if reconfiguring is required, you need to find out when is the best time to perform this step: before or after the migration of the operating system. In Chapter 8, “Applications” on page 179, we give some detailed information on this task for some of the most used applications, such as SAP, and for databases, such as DB2® and Oracle.

4.4.1 Outline of possible migration scenarios

In the following chapters, we give some example test scenarios of various migration methods. These may be similar to your environment, which means you can directly use the information. Your environment might be more complex or might vary slightly in some of the aspects, but the principle of the installation method and their prerequisites should still be the same. Outlining several possible scenarios should, on one hand, help you with your migration planning and preparation. On the other hand, using the specific examples in Chapter 5, “Migration by media” on page 79 and Chapter 6, “Migration by NIM” on page 105 should help you run into as few problems as possible when migrating your production environment.

Migration of stand-alone machine using media

The first case which we consider in Chapter 5, “Migration by media” on page 79 is a migration of an AIX Version 4.3 machine to AIX 5L Version 5.2 using media. We describe the migration steps, including available install options and limitations. The setup is shown in Figure 4-1 on page 51.

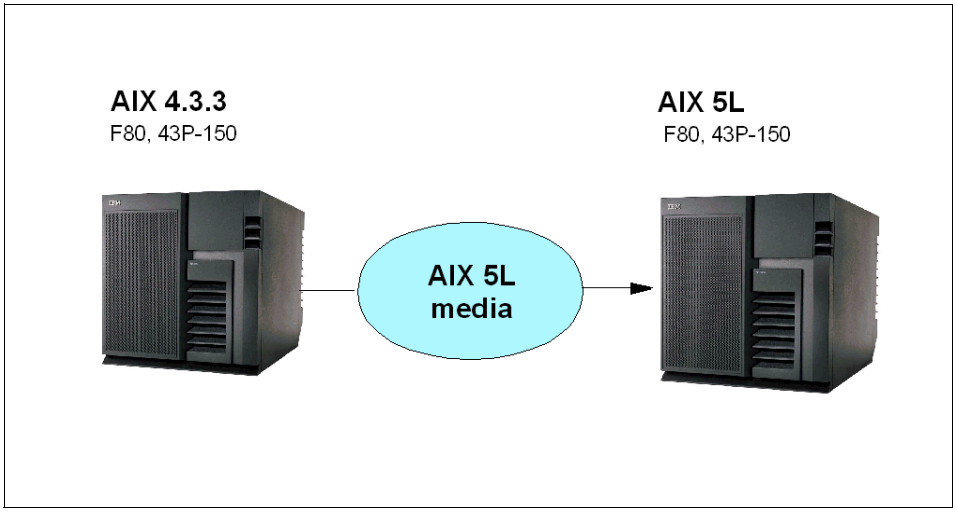

Figure 4-1 AIX migration by media

Furthermore, we discuss the case that you have a POWER3 server running AIX Version 4.3, which you would like to move to a partition of a POWER4 machine, such as a pSeries 690 in our example. POWER4 machines require AIX 5L Version 5.1 and higher, which means you have to migrate your POWER3 machine first to AIX 5L. The next step is to make a backup by tape, CD, or DVD and restore the tape or CD/DVD onto the p690. We include the various backup options available with AIX 5L Version 5.1 and 5.2 and outline which method is the most practical and reliable one. The setup is shown in Figure 4-2 on page 52.

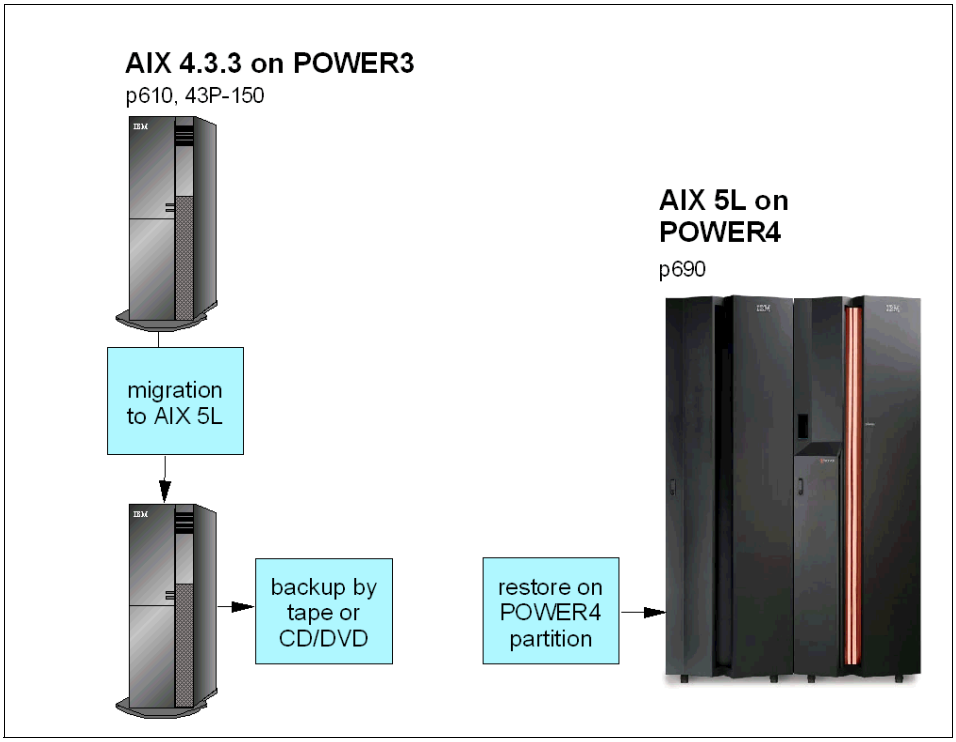

Figure 4-2 Migration of AIX 4.3 POWER3 machine to AIX 5L POWER4 partition

Migration of NIM master

In the next example in 6.3, “NIM master migration” on page 117, we are going to discuss the migration of the AIX Version 4.3 NIM master machine. To migrate this machine to AIX 5L Version 5.2, you need to use the media described in Chapter 5, “Migration by media” on page 79. There are different approaches possible. You can either directly migrate your existing NIM master, which has the disadvantage that it is unavailable during the migration process. Or, alternatively, you can clone your NIM master machine onto another server, migrate this machine, and set it up as NIM master for AIX 5L. The advantage is that at all times you have your AIX Version 4.3 NIM master running. Additionally, we include the steps that are necessary prior to the migration, and the configuration steps that are necessary after the migration to configure the AIX machine as NIM master again.

The second scenario we concentrate on in 6.3, “NIM master migration” on page 117 has to do with moving your applications from a POWER3 machine running AIX Version 4.3 to a POWER4 machine running AIX 5L, with the addition that the POWER3 machine is the NIM master of your environment. We outline

three different approaches for this scenario. The first method is to restore the NIM database of the AIX Version 4.3 system onto a newly installed partition of the p690. The second option is to migrate the POWER3 machine to AIX 5L Version 5.2, back up the NIM database, and restore it onto a newly installed partition of p690. The third way for this scenario is to migrate the AIX Version 4.3 system to AIX 5L, make a backup of the system data, and restore it on the POWER4 machine. A possible example setup is similar to Figure 4-3.

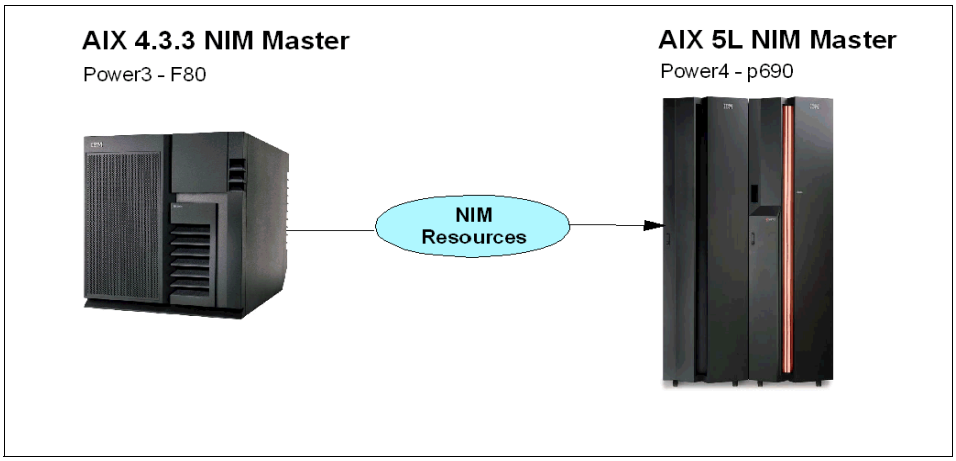

Figure 4-3 Migration scenarios for NIM master

Migration of NIM client

The other important issue that comes with migration is the migration of your NIM clients. This is described, with some test examples, in 6.4, “NIM client migration” on page 137. There are three different possibilities to move your AIX Version 4.3 NIM client to AIX 5L. These are as follows:

- NIM new and complete overwrite
- Alternate disk migration
- NIM migration

We concentrate on the last two methods, which are the most interesting and most useful ones for customers. For this setup, you need your AIX Version 4.3 NIM client as well as a NIM master that runs AIX 5L connected to the same network. A graphical example setup of the alternate disk migration is shown in Figure 4-4 on page 54.

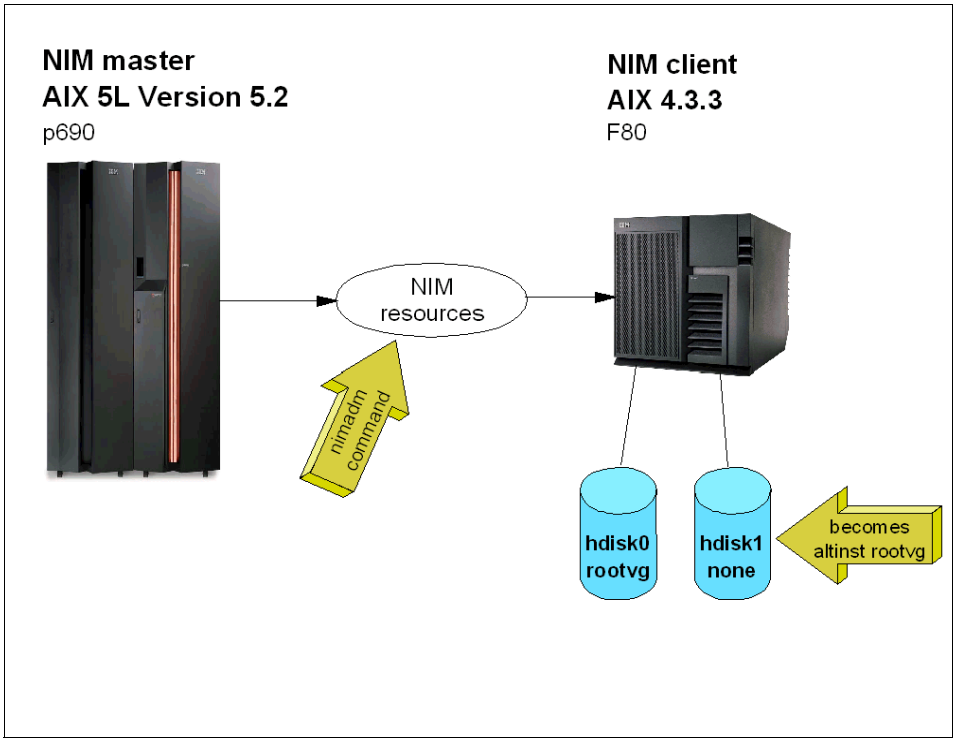

Figure 4-4 Alternate disk migration

4.4.2 Install options

When you plan your migration from AIX Version 4.3 to AIX 5L, you need to make some decisions regarding installation options. One question, for example, is whether you want to have JFS or JFS2 file systems in your migrated system. Another issue is the choice between the 32-bit and 64-bit kernel, as well as whether or not you want the Trusted Computing Base (TCB) enabled. Let us look at some of the possible choices, including their advantages and drawbacks.

There is some limitation with the Trusted Computing Base (TCB) when doing the alternate disk migration. If your AIX Version 4.3 system has TCB enabled and you want this feature enabled with AIX 5L as well, you are required to do a migration by media or Network Installation Management (NIM), as we describe in Chapter 5, “Migration by media” on page 79 and Chapter 6, “Migration by NIM” on page 105. Keeping the TCB enabled is not a possible option when you decide to migrate your system using alternate disk migration. The reason for this problem and workarounds for it are described in 6.4.3, “Alternate disk migration

of client” on page 140. If you do not want TCB enabled with AIX 5L you can use all three methods to migrate your system.

AIX 5L offers to run with the 32-bit or with the 64-bit kernel version. This means that you can choose which kernel to run at installation time or change it after the migration process by relinking it. Both kernels are able to run 32-bit and 64-bit applications. If you decide to enable the 64-bit kernel, it will be linked so that it becomes the running kernel on the system. However, when you perform a migration, you are not able to specify the 64-bit kernel. This is only possible if you do a new and complete overwrite installation. Nethertheless, it is not a problem, because as just described, you can relink to the 64-bit kernel once your system is running on AIX 5L.

For new and complete overwrite installations of AIX 5L Version 5.1, the option of choosing 64-bit kernel enablement automatically installs the JFS2 file system. For a migration installation, however, the option to have the JFS2 file system is not available either. This means that with a migration installation, you cannot have JFS2 file systems. The only way to migrate your JFS to JFS2 is by doing it manually. There are two ways of performing this step. If you have enough space available in the volume group, you can create a new JFS2 file system with the same attributes and copy the files. The alternative is to make a backup of the data in your existing file system first. Then create a new file system with an Enhanced Journaled File System (JFS2) and restore the backup in the JFS2.

In summary, when you perform a migration installation on your AIX system, the kernel and the file systems will remain as they are. For the kernel, this can be changed at a later time using a relink to the other kernel. Migrating your JFS to JFS2 file system is possible only manually. In addition, Table 1-2 on page 7 displays the functional differences between these two file system options.

4.4.3 Migration schedule

The most frequently asked question is how long does a migration from AIX Version 4.3 to AIX 5L take? The answer is not that simple. It depends on many different factors, such as what hardware is used, what migration method, do we migrate from AIX Version 4.3 to AIX 5L Version 5.1 or 5.2, and many others. This means that we cannot give exact migration times for each machine, but we will provide you with some estimated times for various machines and migration methods that we have tested as part of this migration guide. This should give you a rough idea how much time to schedule for the actual migration process.

There are two different schedules you need to consider, that is, a planning schedule and a migration schedule. We will discuss both in the following sections.

Sample planning schedule

One of the most important tasks is to plan and prepare your migration carefully. This includes writing a detailed planning schedule for your environment. The more complex your server environment is, the more detailed this schedule should be. To give you an idea of which issues you need to consider, we give a sample planning schedule in this section. This may vary, though, depending upon the size, complexity, and number of systems you intend to migrate.

The planning schedule can be worked on as part of your existing workload. It is not very complex, but still requires that you check all requirements, as described in Chapter 4, “Planning and preparation” on page 37. This ensures that the migration is achievable and that you have all hardware or software components available when the actual migration process starts.

In many cases, testing and porting third-party and in-house applications will need a separate schedule. This may even be the most complex and time consuming part of the migration preparation. If you have large and complex environments as well as aggressive deadlines and demands on the downtime of your system, we recommend you set up a test environment first. Verify with your test environment that your planning and migration schedule is realistic and that your applications work as desired on AIX 5L. This ensures that you run into as few problems as possible with the migration of your production environment. The following activities are a guideline for your planning and preparation schedule:

- Documenting your existing environment
- Mapping your existing environment
- Determining if migration is possible
- Determining the migration strategy for each machine of the environment
- Ordering, delivering, and installing new hardware and software
- Setting up the test environment
  - Copy the target system onto the test system.
  - Migrate the test system.
  - Test and port the third-party applications on the test system.
  - Test and port the in-house applications on the test system.
  - Test and port the locally written shell scripts.
  - Make a detailed list of migration and configuration steps done on the test system.
- Documenting all the necessary information for the migration of the production environment, which includes all findings and necessary tasks experienced with the test environment

- Schedule downtime for the migration process and inform users about the downtime schedule

Sample migration schedule

This is the time where you have to schedule your system downtime, as well as downtime for backup purposes before and after the migration. If you take your system out of production to make significant changes to it, it is *essential* to have a working backup of your AIX Version 4.3 machine. We describe this topic in detail in 4.6, “Backup of existing environment” on page 65.

There are additional post-migration steps that you must do. These include getting (again) a working backup of your newly migrated system as well as getting your applications running. The amount of time and work you have to do at this stage depends on how well you have tested your applications prior to the migration on AIX 5L. You may have to upgrade your applications by applying fixes or even install a new version of the application. We also recommend you run the **snap** command again right after the migration process to have a comparable system configuration. We give detailed information on post migration issues in Chapter 7, “Post migration tasks” on page 165.

Furthermore, what you are most interested in is a time frame for a migration installation. The duration of the migration varies with different migration methods, and also depends strongly on the hardware you use. With an RS/6000 Model 43P-150 with 1 GB of RAM, which we use in our test environment it takes, for example, about an hour to migrate to AIX 5L Version 5.1. The alternate disk migration using an IBM @server pSeries F80 with 2.6 GB data and 2 GB paging space as the client machine and an IBM @server pSeries F80 as the NIM master takes about two hours. However, these numbers only give you a rough idea, but might vary considerably in your environment, with a lot of data on the servers, and also depends on whether you migrate to AIX 5L Version 5.1 or 5.2. The other important issue is the fact that the preparation and planning process will probably be the more time-consuming task and the migration process itself is done within hours.

4.5 Documentation of existing environment

There are various reasons why we recommend that you document your environment and system configuration prior to starting the migration process. It guarantees you a working environment with the same or enhanced functionality after the migration to AIX 5L. Furthermore, it gives you a good overview with detailed information on your existing environment and helps you to determine whether or not the migration installation is possible at all. Another reason is that

this documentation enables you to create certain NIM groups of similar machines for which you can create customized configuration files for future use.

The amount of work involved in getting your machines and your server environment documented depends strongly on the number and the complexity of your servers. If you have, for example, many similar servers that you keep synchronized, you only need to do the documentation once. If you have, on the other hand, many different machines with various hardware and software components, you need to examine each machine individually. This means that when you plan your migration schedule, allow time for getting your environment documented properly.

There are several ways you can get your machines documented. One method is to use the **snap** command. This will be explained in detail, including some example output, in 4.5.1, “Using standard AIX and the snap command” on page 58. Another possibility is to use a variety of AIX commands of your choice, which you could combine using a self-written script in order to collect the information that is important for your documentation. There is a detailed section in this chapter on that topic as well.

4.5.1 Using standard AIX and the snap command

Using a **snap** command, you can obtain the most important system information. It was originally designed as a problem determination tool for support staff to assist customers with problems when troubleshooting their system. Almost all system-related information is captured and stored in the /tmp/ibmsupt directory. It is a very valuable source of information. We recommend that you run the **snap** command before and after the migration and check to see if there are any differences in the configuration that might be of any concern to your production environment.

Even though the information collected is very useful, you might not like the format of the output files. Therefore, our recommendation is to run the **snap** command prior to starting your migration progress. If you feel fine about the output, go ahead and use it to document your AIX system. If you rather write your own script, you may still like to collect some of the main important commands that run using **snap**, as shown in Example 4-5 on page 60 and in Example 4-6 on page 60.

To capture all system information, enter the **snap -a** command as root. The command runs for several minutes, depending on the complexity and performance of your machine. There is temporary disk space of about 8 MB required to execute all commands. There are several different options that enable you to gather only some of the information. For detailed information on

possible flags and options, refer to the *AIX Version 4.3 Commands Reference, Volume 5*, SC23-4119, which can be found at:

http://publib.boulder.ibm.com/cgi-bin/ds_form?lang=en_US

We will not show you an example of a complete system configuration collected using the **snap** command, because it is too detailed. However, we want to give you some idea of the output files. We have included two of them in Appendix A, "Hardware and fileset support with AIX 5L" on page 197. These include all the information you see in the filesys.snap and tcpip.snap files. As mentioned before, you find the output directories of all gathered information by running **snap** in /tmp/ibmsupt, if not specified differently. This is shown below in Example 4-4.

Example 4-4 Output directories after running the snap command

```
# ls -al /tmp/ibmsupt
drwx------  19 root     system       512 Jan 29 16:33 .
drwxrwxrwt   8 bin      bin         1024 Feb 04 12:40 ..
drwx------   2 root     system       512 Jan 29 16:35 XS25
drwx------   2 root     system       512 Jan 29 16:35 async
drwx------   2 root     system       512 Jan 29 16:35 dump
drwx------   2 root     system       512 Jan 29 16:35 filesys
drwx------   2 root     system      1024 Jan 29 16:33 general
drwx------   2 root     system       512 Jan 29 16:35 install
drwx------   2 root     system       512 Jan 29 16:35 kernel
drwx------   2 root     system       512 Jan 29 16:35 lang
drwx------   2 root     system       512 Jan 29 16:35 lvm
drwx------   2 root     system       512 Jan 29 16:35 nfs
drwx------   2 root     system       512 Jan 29 16:33 other
drwx------   2 root     system       512 Jan 29 16:35 printer
drwx------   2 root     system       512 Jan 29 16:35 sna
drwx------   2 root     system       512 Jan 29 16:35 ssa
drwx------   2 root     system       512 Jan 29 16:35 tcpip
drwx------   2 root     system       512 Jan 29 16:33 testcase
drwx------   2 root     system       512 Jan 29 16:33 wlm
```

We choose to look at the filesys directory as an example. In this directory, you find a filesys.snap file, which contains all information to do with file systems. A collection of the commands used to obtain this information is displayed in Example 4-5 on page 60. It gives you an idea of the amount of information gathered. For the complete output of the filesys.snap file, refer to Appendix A, "Hardware and fileset support with AIX 5L" on page 197. Some of the commands shown below will be explained in more detail in 4.5.2, "Documenting your hardware" on page 61. Additional Logical Volume Manager (LVM) information collected by the **snap** command can be found at /tmp/ibmsupt/lvm/lvm.snap on your machine after you run the **snap** command.

Example 4-5 Collection of commands used to obtain the filesys.snap file

```
df -k
mount
lsfs -l
lsvg
lsvg -p rootvg
lsvg | xargs lsvg -l
lspv -l hdisk0
lspv -l hdisk1
lslv -l hd5
lslv -l hd6
lslv -l hd9var
lsattr -El hdisk0
lsattr -El hdisk1
```

We show another sample collection of commands in Example 4-6. We now collect all network related information in the tcpip directory. You find the gathered information in the tcpip.snap output file. We show, in this section only, the commands used rather than the complete output. For the complete output, refer to Appendix A, “Hardware and fileset support with AIX 5L” on page 197.

Example 4-6 Collection of commands used to obtain the tcpip.snap file

```
lssrc -a
netstat -m
netstat -in
netstat -v
netstat -s
netstat -an
netstat -sr
netstat -nr
no -a
```

As you can see, the commands used as part of the **snap** analysis are well-known AIX commands. Many of these commands begin with the letters ‘ls’ for listing system information. It is a nice collection of commands that you can either select to write your own script with or which you can run as they are using the **snap** command to document your environment. The easiest and quickest way is to use the pre setup **snap** command. This gives you confidence knowing for sure that all needed important information is selected prior to your operating system migration.

4.5.2 Documenting your hardware

In 4.1, “Hardware requirements” on page 38, we provide information on supported and unsupported hardware for AIX 5L. You need to collect and document all hardware components of your AIX Version 4.3 system and use this information to determine if each component is still supported with AIX 5L. This applies to IBM hardware as well as to third-party components. In order to have a functioning system after migration, you need to make sure that all installed and attached hardware components will work and are supported after the migration to AIX 5L.

We give some examples in this section that are complementary to 4.5.1, “Using standard AIX and the snap command” on page 58. Looking at the installed memory, for example, use the commands shown in Example 4-7. This information is collected automatically at /tmp/ibmsupt/general/general.snap when running the **snap** command.

Example 4-7 Installed memory information

```
# lsdev -CHc memory
# bootinfo -r
```

Similarly, you can obtain information on disk drives using the **lsdev** command. This way, you collect descriptive information of the device as well as information on location and status. We show a sample output in Example 4-8.

Example 4-8 Disk drive information using the lsdev command

```
# lsdev -CHc disk
name   status    location     description
hdisk0 Available 11-08-00-2,0 16 Bit LVD SCSI Disk Drive
hdisk1 Available 11-08-00-4,0 16 Bit LVD SCSI Disk Drive
hdisk2 Available 27-08-00-8,0 16 Bit LVD SCSI Disk Drive
hdisk3 Available 27-08-00-9,0 16 Bit LVD SCSI Disk Drive
hdisk4 Available 27-09-00-8,0 16 Bit LVD SCSI Disk Drive
```

The location code gives information on the slot number of the controller and the address of the device. For a detailed description of the format and meaning of the location code, refer to your system documentation. In our case, the slot number is the second part of the location number, that is, 08 and 09. The last part of the location number indicates the SCSI address of the device. For hdisk0 and hdisk1, for example, that is 2 and 4 respectively.

To extract information about other non-disk SCSI devices, use the **lsdev -CHs scsi** command. This will display information about your tape drive or CD-ROM. Similarly, you can get adapter information as displayed in Example 4-9 on page 62.

Example 4-9 Extracting adapter information

```
# lsdev -CHc adapter
name    status    location description

sa0     Available 01-S1    Standard I/O Serial Port
sa1     Available 01-S2    Standard I/O Serial Port
sa2     Available 01-S3    Standard I/O Serial Port
sa3     Available 01-S4    Standard I/O Serial Port
siokma0 Available 01-K1    Keyboard/Mouse Adapter
fda0    Available 01-D1    Standard I/O Diskette Adapter
scsi0   Available 11-08    Wide/Ultra-2 SCSI I/O Controller
scsi1   Available 11-09    Wide/Ultra-2 SCSI I/O Controller
scsi2   Available 27-08    Wide/Ultra-2 SCSI I/O Controller
scsi3   Available 27-09    Wide/Ultra-2 SCSI I/O Controller
mg20    Available 2D-08    GXT130P Graphics Adapter
sioka0  Available 01-K1-00 Keyboard Adapter
ppa0    Available 01-R1    CHRP IEEE1284 (ECP) Parallel Port Adapter
ent0    Available 17-08    IBM 10/100/1000 Base-T Ethernet PCI Adapter
(14100401)
tok0    Available 1A-08    IBM PCI Tokenring Adapter (14103e00)
ent1    Available 21-08    IBM 10/100 Mbps Ethernet PCI Adapter (23100020)
ent2    Available 3A-08    IBM 10/100/1000 Base-T Ethernet PCI Adapter
(14100401)
sioma0  Available 01-K1-01 Mouse Adapter
```

The information obtained in Example 4-8 on page 61 gives you the physical structure of the disks. Furthermore, it is important to gather information on volume groups, logical volumes, and file systems. For a list of configured volume groups, use the **lsvg** command. To extract information from the root volume group, for example, use the **lsvg -p rootvg** command. For a list of configured file systems, type **lsfs**; for file system size information, see Example 4-10.

Example 4-10 Output of the df -k command

```
# df -k
Filesystem    1024-blocks      Free %Used    Iused %Iused Mounted on
/dev/hd4            16384      8780   47%     1098    14% /
/dev/hd2           507904     44336   92%    15365    13% /usr
/dev/hd9var         16384     14440   12%      155     4% /var
/dev/hd3           114688    106428    8%      112     1% /tmp
/dev/hd1            16384     15820    4%       20     1% /home
```

4.5.3 Documenting your software

After you have a complete documentation of your hardware either using the well-suited **snap** command or your own method of getting the most important

system information, you need to gather installed software information as well. This should exist in a production environment in great detail. However, from experience, this is not true for many customers. Therefore, we recommend gathering all this information and documenting it appropriately before starting your migration of the AIX operating system. This serves two purposes: You have a complete list of software that is needed for the migration process and which is additionally of great use for future reference beyond the migration process.

Base operating system

You should document the level of the base operating system that runs on your AIX machines. This can be done by using the **oslevel** command. If you use the -r flag, you can also obtain information on which maintenance level your system is running. In Example 4-11, we show an example output.

Example 4-11 Operating system level

```
# oslevel
4.3.3.0

# oslevel -r
4330-08
```

As you can see from the output, we have AIX Version 4.3.3 Maintenance Level 8 (ML 08) installed.

The next step is to document the installed Licensed Program Products (LPPs) of your AIX machine. For this task, you can use the **lslpp -l** command or the **lslpp** command with a variety of other flags, depending on how much detail on the fileset you want. If you use the **snap** command to document your environment, you do not need to document your operating system. The **snap** command documents the list of installed LPPs in the /tmp/ibmsupt/general/general.snap file. The way the information is gathered there is with the **lslpp -lc** and **lslpp -La** commands.

AIX license

Another important issue is the number of licensed users associated with your AIX machine. To check this number on your system, run the **smitty chlicense** fast path and you will get a similar output to what is shown in Example 4-12 on page 64. The number of licensed users in our case is 40.

Example 4-12 AIX user license using SMIT

```
                Change / Show Number of Licensed Users

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

                                                        [Entry Fields]
Maximum number of FIXED licenses                        [40]
FLOATING licensing                                      off


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

You should have license information associated with your IBM @server pSeries or RS/6000 machine as well. Furthermore, you can gather the same information we obtained in Example 4-12 by typing the **lslicense** command on the command line. Example 4-13 shows the output of the command.

Example 4-13 AIX user license using the command line

```
# lslicense
Maximum number of fixed licenses is 40.
Floating licensing is disabled.
```

4.5.4 Additional documentation

Another very useful command to get the most important parameters for the AIX operating system is **lsattr -El sys0**. The output displays parameters for the sys0 device. Similarly, almost all other devices can be specified and documented by using the **lsdev**, **lsattr**, or **lscfg** commands. An example output is shown in Example 4-14.

Example 4-14 Parameters for the AIX operating system

```
# lsattr -El sys0
keylock       normal        State of system keylock at boot time         False
maxbuf        20            Maximum number of pages in block I/O         True
maxmbuf       0             Maximum Kbytes of real memory allowed        True
maxuproc      128           Maximum number of PROCESSES allowed          True
autorestart   false         Automatically REBOOT system after a          True
iostat        false         Continuously maintain DISK I/O history       True
realmem       393216        Amount of usable physical memory in          False
conslogin     enable        System Console Login                         False
fwversion     IBM,TCP02007  Firmware version and revision                False
maxpout       0             HIGH water mark for pending write            True
```

```
minpout       0             LOW water mark for pending write            True
fullcore      false         Enable full CORE dump                       True
pre430core    false         Use pre-430 style CORE dump                 True
rtasversion   1             Open Firmware RTAS version                  False
modelname     IBM,7043-150  Machine name                                False
systemid      IBM,011090277 Hardware system identifier                  False
boottype      disk          N/A                                         False
SW_dist_intr  false         Enable SW distribution of interrupts        True
```

To gather information about third-party software installed on your system, you may be able to get some information when you run the **lslpp** command. There is no consistent method, however, to get a complete list of applications and their version number that are running on your system. In most cases, you see the version and release number of the third-party application when you start it. Similar arguments apply for in-house developments, including customized shell scripts. It is the customers responsibility to document changes and updates. To make sure the in-house applications run with AIX 5L, you need to test them prior to the migration to AIX 5L.

Configuration files for users, groups (including their attributes), and passwords will be preserved during the migration. For completeness, you may want to include this information in your system documentation. As with all the examples before, if you use the **snap** command as recommended, you automatically get this information documented. The file is at /tmp/ibmsupt/general. Furthermore, it is your responsibility to document any additional personal configurations, such as desktop settings, any other configurations where the values differ from the default values, and to save important information that you store in /tmp prior to the migration of AIX. As explained in 4.3.2, “Migration installation” on page 46, almost all directories, such as /home, /var, /usr, the root volume group, logical volumes, and your system configuration and previously installed software, is preserved. The only file system that will be new after the migration is /tmp.

4.6 Backup of existing environment

Backing up your existing environment is necessary to ensure that none of your system data gets lost. This way, you protect your system against any problems that may occur during or after the migration process. In general, our experience with migrating systems from AIX Version 4.3 to AIX 5L have been very positive. However, having an adequate backup of your system means you can easily restore your environment to its original state if needed.

4.6.1 Backup options

There are two different ways to back up large amounts of data: The complete system backup and the incremental backup. We will concentrate on the complete system backup in this chapter because that is what you are going to need before starting your migration from AIX Version 4.3 to AIX 5L. However, an incremental backup is useful for storing data that changes often on the system. It allows the administrator to make sets of incremental backups each night of the week, for example, and to do a full backup only once every weekend.

There are, in addition, different ways of backing up your user or application data and creating an mksysb image of your machine. These will be explained in detail in the next sections. However, before making a decision on your backup strategy, you need to understand file system structure, data placement and how each of these methods work.

A file system is a section of your hard disk that has been allocated to contain files. It is accessed by mounting the file system over a directory. To the user, the mounted file system looks just like any other directory. There are structural differences between the file systems and directories, though, and the data within these can be managed separately. The first time you install the AIX operating system on your machine, it is loaded into a directory structure in the /(root) file system tree.

Additionally, there are two different types of data. The system data makes up the operating system and its extensions and is kept in the system file systems, such as /(root), /usr, /tmp, or /var. Furthermore, there is user data that is kept in the /home file system or any other user-specific file systems. You ought to manage your user and system data separately for backup and restore reasons. This means you back up your system data in a different image to your user data. In addition, system data does not change as often as user data. A separate structure ensures that this backup image will be smaller. Similar reasons apply for restoring user and system data. Restoring only user data saves time and effort because a reboot, as with restoring system data, is not necessary.

Your backup strategy before doing the migration should therefore be to unmount all user file systems, including /home, with the **umount** command. For this task, the file systems must not be in use. The only mounted file systems should be /(root), /var, /usr, and /tmp. Then back up your operating system data using the **mksysb** command. After that is finished, mount the user file system, back up files, file systems, or other non-root volume groups.

There are different commands to be used for saving your user and system data. Regardless of whether you use the Web-based System Manager, SMIT, or the command line, AIX uses four commands to create backups. These are **mksysb**, **savevg**, **mkcd**, and **restvg**. The **mksysb** command is used to clone the root volume

group. The **savevg** command is used for user volume group backups and the **restvg** is used to restore the user volume group. All of these will be explained in more detail further below in 4.6.2, “Backup of root volume group” on page 69 and in 4.6.3, “Backup of a non-root volume group” on page 72.

The **mkcd** command backs up a volume group (either the root volume group or a non-root volume group) to writable CD or DVD media. It transfers a previously made image to the media or calls the **mksysb** or **savevg** command to create the image that will be placed onto the media. An example of the command usage is shown in Example 4-15.

Example 4-15 Use of mkcd command

```
# mkcd -d /dev/cd0
# mkcd -d /dev/cd0 -m /mksysb_image/mksysb5
```

With this command, we back up the operating system of the AIX machine. The first line of Example 4-15 creates a bootable CD on a CD-ROM named /dev/cd0. The second line of the same example uses a previously created mksysb image which we call mksysb5. The -d flag is required and indicates the CD-ROM device.

Furthermore, for backing up your root volume group, you need to decide between a backup with which you can reinstall the same system and a backup with which you can reinstall another system. With AIX 5L Version 5.2, this separation is no longer needed, because at installation or migration time, all device drivers will be installed by default, even if a device is not attached to the machine. This makes it easier to clone a system using the **mksysb** or **mkcd** command. Table 4-3 on page 68 gives an overview of these backup options. You can use Web-based System Manager, SMIT, or command line.

Table 4-3 Making a backup of an AIX Version 4.3 system

Task	Web-based System Manager	SMIT	Command Line
Create a backup of rootvg	► Type **wsm** on the command line. ► Choose the **Backups** container, then **Backup of System**.	► For backup on tape, run **smitty mksysb**. ► For backup on CD/DVD, run **smitty mkcd** (to install the same system) or **smitty mkcdgeneric** (to install on another system).	► Run the **mksysb** command, as described in Example 4-17 on page 71. ► Run the **mkcd** command, as described in Example 4-15 on page 67.
Create a backup of non-rootvg (datavg)	► Type **wsm** on the command line. ► Choose the **Volumes** container, then **datavg** and **Back up**.	► For backup on tape or file, run **smitty savevg**. ► For backup on CD/DVD, run **smitty savevgcd**.	► Run the **savevg** command, as described in Example 4-19 on page 73.

Additional information

There are some issues you need to be aware of when planning your backup strategy. When you have database applications that write directly to the device, you often use raw devices. These do not contain a JFS or JFS2 file system. The image created by the **mksysb** or **savevg** commands does not include data on raw devices or in paging space. Additionally, when you have special features installed or when your system uses device drivers that are not shipped with the operating system, these device configurations are not backed up using the two AIX backup commands.

Furthermore, you have the possibility to exclude certain files from the backup. For this task, you can use the /etc/exclude.rootvg file where you list the appropriate files.

The **mksysb** command requires working space for the duration of the backup. Hence, for a backup of an AIX Version 4.3 system, you need to have enough space in the /tmp file system. Including the -X flag when starting the **mksysb** command, as we show in Example 4-17 on page 71, ensures that the system increases the size of /tmp automatically during the backup if needed. Before you start the backup, we recommend checking the integrity of the file system with the **fsck** command.

4.6.2 Backup of root volume group

The root volume group is a hard disk or group of disks that contains startup files, the operating system, configuration details, and any optional software. It can be backed up using the **mksysb** command, which creates an image onto tape or in a file. If you create the backup on tape, the tape is bootable. It includes the installation programs needed to install from the backup. When you back up your system using CD or DVD, there are two cases you need to consider. These are a backup to reinstall the same system and a backup to reinstall another system. The difference and when to use which method is explained in the following sections. This differentiation is not necessary for backups of your AIX 5L Version 5.2 system anymore. By default, all device drivers are installed, which means you can restore the backup on any machine. However, for a backup of your AIX 5L Version 5.1 system, it is a concern, so we explain it in the following sections.

Backup to reinstall the same system

This backup option is useful if you want to use the same machine after the migration and only upgrade the operating system. Using the **mksysb** command, you back up the operating system or root volume group in a form that is bootable only on the source system. To reinstall it on another system, you need the original media to retrieve the different filesets. This is the method most customers use to have a copy of their data before starting a migration installation. When you use a CD or DVD instead of a tape, this backup method is called a *personal backup*.

Backup to reinstall another system

The second option is useful when you want to bring your applications from a machine running AIX Version 4.3 onto a POWER4 machine running AIX 5L, as described in 4.3.4, “The mksysb installation” on page 48. Another use would be in an environment where each machine needs to have the same image installed, but has different hardware configurations. Again, you back up your root volume group on your migrated AIX 5L Version 5.1 system, but in such a way that it can be booted and restored onto another system without the original AIX media. This type of backup can be made to tape and file to be used by NIM to clone the operating system environment onto multiple remote machines as well. It can additionally be made to CD or DVD and in this case it is called a *generic backup*.

As indicated in Table 4-3 on page 68, you can use either the Web-based System Manager, SMIT, or command line to start a backup of your system. Here is an example of how to make a backup of the rootvg to a tape device named /dev/rmt0, first using the Web-based System Manager, as shown in Figure 4-5 on page 70. To get to this screen, type **wsm** on the command line, then click on the **Backups** container and finally click on **Back up System**.

.

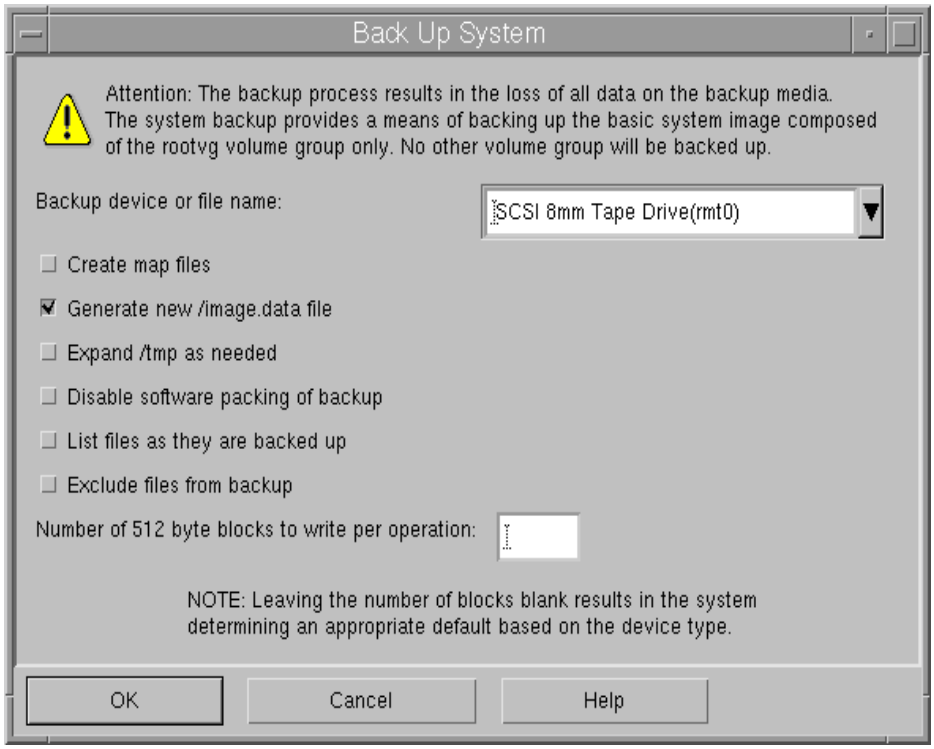

Figure 4-5 Backup menu using the Web-based System Manager

Secondly, we show the same example using SMIT and get the screen shown below in Example 4-16 by typing the fast path **`smitty mksysb`**.

Example 4-16 Backup menu using SMIT

```
                        Back Up the System

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

                                                        [Entry Fields]

WARNING:  Execution of the mksysb command will
          result in the loss of all material
          previously stored on the selected
          output medium. This command backs
          up only rootvg volume group.

Backup DEVICE or FILE                                   []
Create MAP files?                                       no
EXCLUDE files?                                          no
```

```
List files as they are backed up?                          no
Generate new /image.data file?                             yes
EXPAND /tmp if needed?                                     no
Disable software packing of backup?                        no
Number of BLOCKS to write in a single output               []

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

The third way to start a system backup is by using the command line. This is shown in Example 4-17, first, as before, to a tape and then to a file named mksysb5. The -i option creates a new image.data file. This contains information about the sizes, names, maps, and mount points of all the file systems and logical volumes in your root volume group. The -X flag automatically expands the /tmp file system to whatever size the **mksysb** command requires. This way, the backup is not aborted if the backup runs out of space in /tmp.

Example 4-17 Use of the mksysb command

```
# mksysb -X -i /dev/rmt0
# mksysb -X -i /mksysb_image/mksysb5
```

To make sure your backup is complete, remember that only mounted file systems within the rootvg are archived with the **mksysb** command. It does not back up file systems mounted across an NFS network or any other unmounted journaled file systems. You must unmount any local directory that is mounted over another local directory; otherwise, the backup will create two physical copies of the mounted directory. This duplication might exceed a threshold or number of files allowed in file system. It could lead to backup errors, which cause future installations of the backup image to fail.

Another important issue is to have all external hardware installed, including external devices such as tape and CD-ROM drives before you start the backup to avoid inconsistency in data. This is not necessary anymore for AIX 5L Version 5.2 systems because all device drivers will be installed by default, even for non-attached devices.

After the restore of your backup, the same root and user passwords will be active. This might be a security concern in your environment and an issue to remember.

4.6.3 Backup of a non-root volume group

In addition to backing up your system data, you need to make sure your user data is protected and saved before starting the migration installation. With the **savevg** command, you can back up other non-root or user volume groups where you keep your user data files and application software. The non-root volume group must be varied on and the file systems must be mounted. The **savevg** command creates an image onto tape or in a file. It does not create a bootable tape even if the volume group is the rootvg.

The user volume group can be restored with the **restvg** command. An example of how to back up the user volume group called datavg onto tape is shown in Figure 4-6. We started the Web-based System Manager by typing **wsm** on the command line. Then we click on the **Volumes** container. Select the user volume group **datavg** and click on **Back up**. The window in Figure 4-6 should appear.

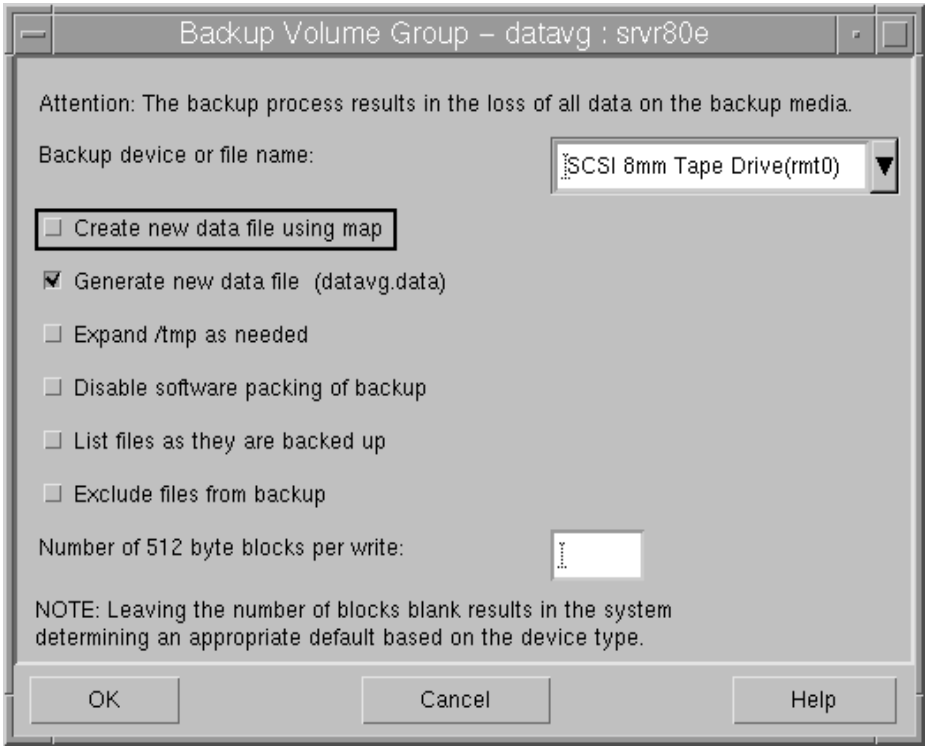

Figure 4-6 Backup of a non-root volume group with Web-based System Manager

This task works in a similar way when using SMIT and the command line. An example of the SMIT screen is shown in Example 4-18 on page 73. The fast path to use is **smitty savevg**.

Example 4-18 Backup of non-root volume group with SMIT

```
                  Back Up a Volume Group to Tape/File
Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                              [Entry Fields]
WARNING:  Execution of the savevg command will
          result in the loss of all material
          previously stored on the selected
          output medium.

Backup DEVICE or FILE                                         []
VOLUME GROUP to back up                                       []
List files as they are backed up?                             no
Generate new vg.data file?                                    yes
Create MAP files?                                             no
EXCLUDE files?                                                no
EXPAND /tmp if needed?                                        no
Disable software packing of backup?                           no
Number of BLOCKS to write in a single output                  []


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

This way, you can exclude certain files and expand the /tmp automatically if needed, as mentioned in Example 4-16 on page 70. Finally, we show the command line for a backup of user data. The -i flag creates a new datavg.data file, similar to what we have seen in Example 4-17 on page 71. This file contains all information about the sizes of all the file systems and logical volumes in the user volume group. The -f flag specifies the device; in our case, it is the tape device /dev/rmt0.

Example 4-19 Back up of non-root volume group using the command line

```
# savevg -i -f /dev/rmt0 datavg
```

4.6.4 Verification and restore of backup

In order to verify your backup, you can do some verification checks. It is not a bit-by-bit comparison, but an overall check of the contents of the backup. The content list verifies most of the information on the tape, but does not verify that the tape can be booted for installations. You can only know for sure that you can boot from tape if you actually boot from the tape.

For the backup verification process of your rootvg using the Web-based System Manager, select the **Backups** container, click **Selected** and **View Backup Contents**. Using the **smitty lsmksysb** command to list all the files in a system image, you should get a screen similar to Example 4-20. Enter the file or device for the backup and verify the block number and block size.

Example 4-20 Verification of root volume group backup using SMIT

```
                    List Files in a System Image

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

                                                        [Entry Fields]
DEVICE or FILE                                          [/dev/rmt0]
Number of BLOCKS to read in a single input              []
Verify BLOCK size if tape device?                       no


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

To verify a backup using the command line, use the **listvgbackup** command. An example is shown below in Example 4-21. The first line in this example indicates what command to use to list the content of a rootvg backup. The second line is used to list the content of a non-root volume group, which was backed up on tape as well.

Example 4-21 Verifying a backup using the command line

```
# listvgbackup -f /dev/rmt0
# listvgbackup -s -f /dev/rmt0
```

However, the only way to ensure that the mksysb image is correct is by performing a test restoration from the media. That means that if it is a bootable backup, you need to test whether it boots on the intended system.

For user volume groups, the command to use is **restvg -f backup_location**. If you want to minimize the space require, add the -s flag. Example 4-22 on page 75 shows two examples. The first line is for restoring a datavg backup of the user volume group datavg. The second line shows how to restore files from a user volume group backup. The command to use for this option is the **restorevgfiles** command. It restores specific files from a rootvg backup.

Example 4-22 Restoring a Backup using the command line

```
# restvg -s -f /datavg_images/datavg1
# restorevgfiles -s -d /tmp -f /datavg_images/datavg1 ./home/file1 ./home/file2
```

For more detailed information on commands and their flag options, refer to the *System Management Guide for AIX Version 4.3*, SC23-4126. This can be found at:

http://publib.boulder.ibm.com/cgi-bin/ds_form?lang=en_US

4.7 Migration validation

After you have worked through all the sections in this chapter, you should be ready to start planning, preparing, and doing the migration of AIX Version 4.3 to AIX 5L. We include a checklist at this stage that you might find useful to go through again in order to see if you have completed all pre-migration tasks.

4.7.1 Hardware and software requirement check list

We start with hardware and software requirements:

- Are all machines in your environment supported by AIX 5L? For a complete list, see Appendix A, “Hardware and fileset support with AIX 5L” on page 197.
- Are all devices, including adapters, supported with AIX 5L?
- Have you verified the microcode level of your system and its devices?
- Have you ordered and received a delivery of missing hardware and software components?
- Have you installed and configured all hardware and software components?
- Are all external devices, such as the tape, CD, or DVD-ROM drives, physically connected?
- Do you have the AIX 5L installation package available, including the Release Notes?
- Do you have the latest maintenance level for AIX 5L? The latest maintenance level of AIX 5L Version 5.1 is ML 03. It can be obtained from the following URL:

  http://techsupport.services.ibm.com/server/aix.fixsearch51

  You have to search for APAR IY32749, which is the AIX 5100-03 Recommended Maintenance Package.
- Have you checked that all applications run on AIX 5L without licensing problems?

- Do you have enough tapes for doing the system and data backup before and after migration?
- Do you have at least 12 MB of contiguous disk space available for the boot logical volume (hd5)?

4.7.2 Backup and documentation check list

On the documentation and backup issue, go through the following list and make sure you can tick off each item:

- Is your environment documented properly and is this documentation stored in a safe place?
- Is your TCP/IP configuration documented as well?
- Is there any important information in /tmp that you would like to keep? If so, move it and copy it back after the migration process. Otherwise, it will be deleted as part of the migration process.
- Have you got a valid backup copy of the system and user data?
- Have you verified that the backup is restorable to protect yourself from loss of data and functionality?

4.7.3 Additional check list

In addition, there are some more things to verify before starting the migration process. These are:

- Have you decided on an migration method for each machine in your environment?
- Is your time frame and estimated downtime realistic and achievable?
- Have you got a list of ordered migration steps to follow?
- Does the root user have a primary authentication method of SYSTEM? Type **`lsuser -a auth1 root`**. If needed, change the value by typing **`chuser auth1=SYSTEM root`**.
- Have you checked the error report of your machine? If there are problems, solve them prior to starting the migration.
- Have you informed all users about the migration process and an estimated downtime of the system?
- Are any users logged on to the system prior to starting the migration?

4.8 What to do in case migration fails

If you plan and prepare your migration well and take into account the requirements and necessities described in the following chapters, it should be unlikely that you will have a migration failure. When you read through the different examples in the following chapters, you should try to have, in the back of your mind, your own environment and where it differs. The success of your migration will depend strongly on how well you can identify and address possible traps and problems in addition to those mentioned in this book.

If, however, you do experience problems and the migration process is aborted, restore your old backup with AIX Version 4.3 to get your production system running again as quickly as possible. At least you can return your environment to its original state while looking for a solution for your problem.

The next step is to check Chapter 6, “Migration by NIM” on page 105 and the AIX database for known problems. Maybe this problem is already known and there are fixes available. If not, contact IBM to inform them about this issue. If you keep having serious problems with your migration process, maybe you need to consider doing a new and complete overwrite of the operating system and then a transfer of your application data. This should really be the case in a few exceptional environments only and the last solution to use.

5

Migration by media

There are different ways you can migrate your AIX Version 4.3 machine to AIX 5L. The most straight forward method is using media. The process is very similar to a new installation and will be described in detail in this chapter. We first explain the migration steps from AIX Version 4.3 to AIX 5L Version 5.1 with all the available install options and limitations. We then concentrate on what is different with the migration to AIX 5L Version 5.2.

Even though there are some new features in the base operating system installation and the preservation installation, we will not include these two installation methods in our discussion. We concentrate on the migration issues only. When using the migration installation, all logical volumes and file systems on your rootvg are preserved, except for /tmp. Additionally, other volume groups are not affected by the migration process.

For general information on various installation menus available with AIX 5L, refer to the *AIX 5L Version 5.1 Installation Guide and Reference*, SC23-4374 and *AIX 5L Version 5.2 Installation Guide and Reference*, SC23-4389.

The final section outlines the steps to perform when moving from a POWER3 AIX Version 4.3 machine to a POWER4 AIX 5L partition. There are some filesets necessary and tasks to complete to make your AIX system partition ready with AIX 5L Version 5.1. These steps will be discussed in 5.3.2, “Case 1: POWER3 to POWER4 AIX 5L Version 5.1 by tape” on page 89.

Before you start migrating your production systems, make sure you have a validated system backup. For detailed information on pre-migration tasks, refer to Chapter 4, “Planning and preparation” on page 37. Furthermore, we include information on backup by tape and DVD in our test examples in 5.3, “Migration from POWER3 to a POWER4 partition” on page 87.

5.1 Migration to AIX 5L Version 5.1

In this first section we introduce the new features and options that are part of the base operating system migration from AIX Version 4.3 to AIX 5L. We discuss the different advanced options for 32-bit and 64-bit machines as well as the enhanced Journaled File System (JFS2) and possible kernel and desktop choices.

If you insert the AIX 5L base CD into the CD-ROM of your machine and boot the system from CD, you will get the installation menu. For AIX Version 4.3 machines, the default method of installation will be `Migration`, as shown in Example 5-1. As with the new and complete overwrite and the preservation installation, you can define your language environment. That means that you define language and keyboard settings as well as cultural convention.

Example 5-1 Installation menu for AIX 5L Version 5.1

```
                  Installation and Settings

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

    1  System Settings:
         Method of Installation.............Migration
         Disk Where You Want to Install.....hdisk0

    2  Primary Language Environment Settings (AFTER Install):
         Cultural Convention................English (United States)
         Language ..........................English (United States)
         Keyboard ..........................English (United States)
         Keyboard Type......................Default

    3  Advanced Options

>>> 0  Install with the current settings listed above.



88  Help ?          |    WARNING: Base Operating System Installation will
99  Previous Menu   |    destroy or impair recovery of SOME data on the
                    |    destination disk hdisk0.
>>> Choice [0]:
```

In our example, we set all options to English (United States). The third menu, called Advanced Options, will vary depending on whether your AIX machine is a 32-bit or a 64-bit machine. Furthermore, if you do not have a graphics console attached, your desktop choice will be limited.

32-bit systems

When you select the Advance Option menu, there are only two additional choices available. The first option you have on a 32-bit system is the desktop choice. This can be CDE, KDE, Gnome, or NONE. This applies to AIX machines with a graphics console attached. For an ASCII console, you will see that the installation package set is defined with the default value, which means the default configuration bundle will be installed. Furthermore, you can choose to enable or disable the Trusted Computing Base (TCB). This enablement is required for certain security aspects.

64-bit systems

If your AIX server is a 64-bit machine, you have the same advanced options as described for the 32-bit machines. Additionally, for a new and complete overwrite you could choose the 64-bit kernel and enhanced Journaled File System (JFS2). However, when you do a migration installation, you are not given this option. This does not mean that you cannot use the 64-bit kernel on your machine; it is just not an option at migration installation time, but can be enabled once the migration process is finished. The kernel support will automatically be installed. To enable the 64-bit kernel on your AIX 5L system, you need to relink /usr/lib/boot/unix and /unix to point to /usr/lib/boot/unix_64. Then issue the **bosboot** command to write a new boot image and reboot your system. For a detailed discussion of JFS/JFS2 with a migration installation, refer to 4.4.2, “Install options” on page 54.

5.1.1 Migration steps

In this section, we demonstrate an example migration from AIX Version 4.3 to AIX 5 L Version 5.1. We use an RS/6000 7043 Model 150 as the 32-bit test machine and an RS/6000 7025 Model F80 as the 64-bit test machine. We want to use English as the primary language and use the default options in the third menu, the Advanced Option menu.

Step 1. Boot from the AIX media

The first step is to insert an AIX 5L Version 5.1 base CD into the CD-ROM of our machine. We set the bootlist to the CD drive by typing the **bootlist** command, as shown in Example 5-2 on page 83. The first line makes the CD-ROM the first boot device. The second line of the example can be used to display and confirm the configured bootlist.

Example 5-2 Set and verify bootlist

```
# bootlist -m normal cd0

# bootlist -m normal -o
cd0
```

Our AIX server will reboot and display the main installation menu, as shown in Example 5-1 on page 81. Alternatively, you can boot your system into SMS mode and set the boot device to be your CD-ROM drive.

The first screen you will see is similar to Example 5-3, with which you define the system console. You are then required to select a language that will be used during the migration installation procedure. Up to this step, there is no difference from a new and complete overwrite installation of your AIX server.

Example 5-3 Defining the system console

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

Type a F1 and press Enter to use this terminal as the
  system console.
Pour definir ce terminal comme console systeme, appuyez
  sur F1 puis sur Entree.
Taste F1 und anschliessend die Eingabetaste druecken, um
  diese Datenstation als Systemkonsole zu verwenden.
Premere il tasto F1 ed Invio per usare questo terminal
  come console.
Escriba F1 y pulse Intro para utilizar esta terminal como
  consola del sistema.
Escriviu 1 F1 i premeu Intro per utilitzar aquest
  terminal com a consola del sistema.
Digite um F1 e pressione Enter para utilizar este terminal
  como console do sistema.
```

Step 2. Verify the migration installation settings

The next step is to verify the migration installation settings. By default, the installation method will display migration on the disk that you have currently an active version of AIX Version 4.3 installed. You need to make your language and keyboard choices as well as your additional choices, as described at the beginning of this section, and confirm the settings.

Step 3. Begin migration installation

Before the migration process starts, you will see a migration installation summary, which summarizes your selected choices. This is your last chance to

change any settings before the migration of the operating system starts. If the settings are as desired, confirm and the migration process begins.

5.2 Migration to AIX 5L Version 5.2

The steps to migrate your AIX Version 4.3 system to AIX 5L Version 5.2 are exactly the same as for the migration to AIX 5L Version 5.1. The difference is in the available options, which you can choose before starting the migration. We will concentrate on these differences and enhanced installation options of AIX 5L Version 5.2 in the next section.

The first difference is the naming of the installation menu three. With AIX 5L Version 5.1, this is called Advanced Options. In AIX 5L Version 5.2, it is called More Options.

32-bit systems

On a 32-bit machine, you will see the choices displayed in Example 5-4 when you select the **More Options** menu. This means, as with AIX 5L Version 5.1, you have the desktop choice and the Trusted Computed Base (TCB) enablement.

Example 5-4 Migration install options on a 32-bit system

```
                        Install Options

 1.  Desktop..................................................... CDE
 1.  Enable Trusted Computing Base.............................. No
 2.  Import User Volume Groups.................................. Yes
 3.  Enable System Backups to install any system................ Yes
     (Installs all devices and kernels)
 4.  Remove Java 1.1.8 Software................................. No

>>> 0  Install with the current settings listed above.



88  Help ?
99  Previous Menu

>>> Choice [0]:
```

Additionally, you can choose to import user volume groups after the migration process and have all Java 1.1.8 filesets removed automatically. With AIX 5L Version 5.2, by default, all devices and kernels will be installed during the migration of the operating system. The advantage is that if you take a backup of your system, you can restore it on any other machine that supports AIX 5L

Version 5.2. This used to cause problems with earlier versions of AIX, because not all necessary device drivers were on the system backup image.

One limitation of the migration menu is that with an installation, you could choose to install additional software at this stage. This is not an option when you perform a migration installation. However, you can always install additional software after the migration process is finished.

Table 5-1 summarizes the available install options and their default settings when migrating a 32-bit AIX Version 4.3 server to AIX 5L Version 5.2.

Table 5-1 Migration install options

	Available option	**Default option**
Desktop	CDE, GNOME, KDE, or NONE	CDE
Enable Trusted Computing Base	Yes/No	No
Import User Volume Groups	Yes/No	Yes
Enable System Backups to install any system	Yes/No	Yes
Remove Java 1.1.8	Yes/No	No

64-bit systems

For a 64-bit system, the migration settings and advanced options are the same as for the 32-bit system, which means we will not show the same screen shot again. Furthermore, we will not describe the migration steps when going from an AIX Version 4.3 system to AIX 5L Version 5.2. The reason is that they are exactly the same as for the AIX 5L Version 5.1 migration.

We discuss the difference in optional choices in 5.2, “Migration to AIX 5L Version 5.2” on page 84. In addition, we tested this migration process with the same hardware, as described in 5.1, “Migration to AIX 5L Version 5.1” on page 81, which works without any problems.

However, when you migrate to AIX 5L Version 5.2, you see a migration summary before the migration starts. Furthermore, you even get a migration confirmation menu prior to the migration process. This is your last chance to cancel the migration. The menu is shown in Example 5-5 on page 86. Additionally, you can obtain a list of filesets that will be removed and not replaced, as well as other listings that may be of interest.

Example 5-5 Migration confirmation menu

```
                        Migration Confirmation

Either type 0 and press Enter to continue the installation, or type the
number of your choice and press Enter.

    1  List the saved Base System configuration files which will not be
       merged into the system.  These files are saved in /tmp/bos.
    2  List the filesets which will be removed and not replaced.
    3  List directories which will have all current contents removed.
    4  Reboot without migrating.

Acceptance of license agreements is required before using system.
You will be prompted to accept after the system reboots.

>>> 0  Continue with the migration.



88  Help ?
+------------------------------------------------------------------------------
WARNING: Selected files, directories, and filesets (installable options)
from the Base System will be removed.  Choose 2 or 3 for more information.

>>> Choice[0]:
```

An example output of files (which will be removed and not replaced) taken from one of our test migrations is shown in Example 5-6.

Example 5-6 Filesets which will be removed and not replaced

```
bos.powermgt.rte
devices.common.IBM.pmmd_chrp.rte
devices.isa_sio.km.diag
devices.isa_sio.PNP0303.diag
devices.isa_sio.PNP0400.rte
devices.isa_sio.PNP0501.rte
devices.isa_sio.PNP0700.rte
devices.isa_sio.PNP0F03.diag
devices.pci.86808404.com
devices.pci.86808404.rte
devices.pci.PNP0A03.rte
devices.sys.pci.rte
ifor_ls.client.base
ifor_ls.client.gui
pkg_gd
sysmgt.websm.ucf
sysmgt.websm.widgets
```

```
X11.apps.pm
X11.msg.en_US.apps.pm
```

Additionally, some directory content will be removed with the migration to AIX 5L Version 5.2. An example output from our test environment is shown below in Example 5-7.

Example 5-7 Removed directory content

```
/lpp/bos
/tmp
/usr/lpp/bos/bos.rte
/usr/lpp/bos/bos.rte.*
/usr/lpp/bos/deinstl
```

Note: Before you begin the migration process of your production environment, we recommend that you check the output of removed filesets and directory content to avoid loss of data on your systems.

After the migration process is completed, the system reboots as part of the migration procedure. The main migration process is finished. As you can see, it is a very straightforward process. There are some additional tasks to complete after the migration. For these post migration tasks, refer to Chapter 7, “Post migration tasks” on page 165.

5.3 Migration from POWER3 to a POWER4 partition

An important issue at many customer environments is porting their applications from a POWER3 AIX Version 4.3 system to a POWER4 AIX 5L partition. The reason why this is a big issue is that all POWER4 machines that include IBM @server pSeries 630, p650, p655, p670, and p690 are partition enabled but require an AIX level of Version 5.1 or higher.

This means that moving to a POWER4 hardware ensures that you can benefit from the enhanced features of the improved hardware architecture. However, because of the operating system requirements, you need to first migrate your POWER3 machine to AIX 5L, back up the system data, and then restore it onto a POWER4 partition.

One of the issues involved is the fact that you may want to move from a POWER3 server to a partitioned POWER4 machine. To get the operating system partition ready, you need to complete some additional tasks before making the system backup. We will explain these tasks in detail in the following test examples.

Furthermore, we include a detailed discussion on related topics, such as cloning your system to an alternate disk of a different machine as well as backing up your system onto tape or CD/DVD. For this, we guide you through three different test cases that involve the mentioned issues.

5.3.1 Environment setup

Our test environments for the three test cases include the following AIX machines. The setup is displayed in Figure 5-1.

- RS/6000 7043 Model 150: POWER3, 32-bit hardware
- RS/6000 7025 Model F80: POWER3, 64-bit hardware
- IBM @server pSeries 610: POWER3, 64-bit hardware
- IBM @server pSeries 690: POWER4 partition

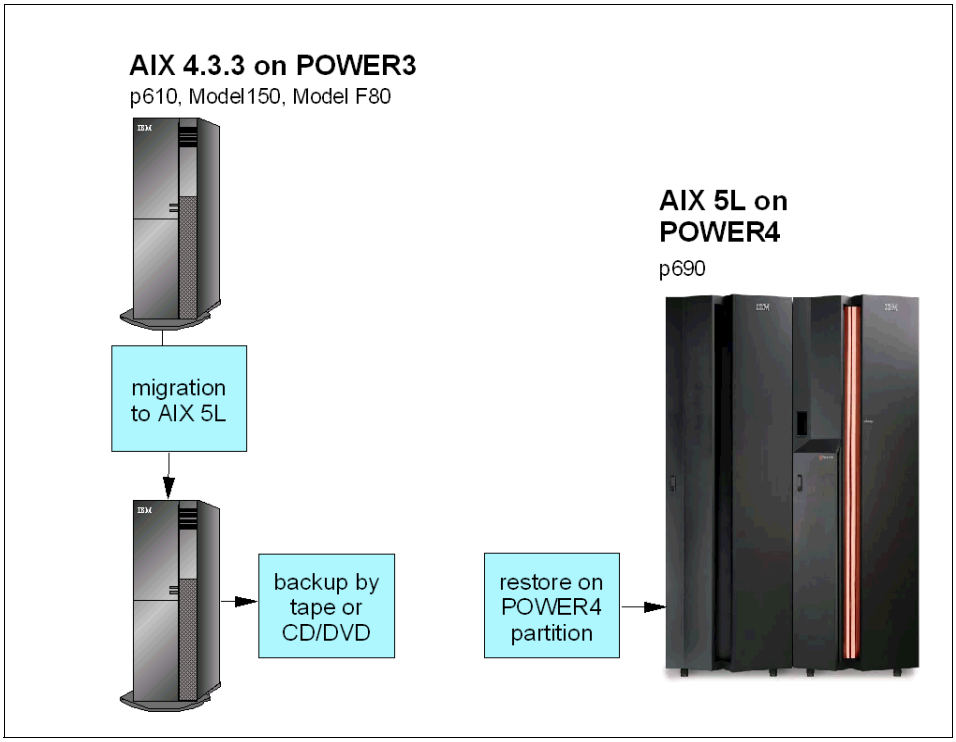

Figure 5-1 Migration of POWER3 machine to AIX 5L POWER4 partition

There are three different cases we will consider. The first case involves an RS/6000 Model 150 as an example of a POWER3 32-bit hardware as well as an RS/6000 Model F80 as an example of a POWER3 64-bit hardware. We migrate both systems to AIX 5L Version 5.1, back up the systems onto tape, and restore them onto a partition of IBM @server pSeries 690.

The second case involves an IBM @server pSeries 610, which is also a POWER3 64-bit machine. We migrate it first to AIX 5L Version 5.2. We then back up the system to DVD and restore it onto a partition of IBM @server pSeries 690.

In the third setup, we clone the rootvg of any of the POWER3 servers to a spare disk, migrate it to AIX 5L, and restore it on an IBM @server pSeries 690 or any other machine. The emphasis in this example is on cloning your system.

- Case 1:
  - Migration of a POWER3 server to AIX 5L Version 5.1
  - Backup by tape
  - Restore to POWER4 partition
- Case 2:
  - Migration of a POWER3 server to AIX 5L Version 5.2
  - Backup by DVD
  - Restore to POWER4 partition
- Case 3:
  - Cloning rootvg to an alternate disk of POWER3 server
  - Migration of cloned disk to AIX 5L
  - Restore on POWER4 partition or any other system

5.3.2 Case 1: POWER3 to POWER4 AIX 5L Version 5.1 by tape

In our first test example, we consider migrating an RS/6000 Model 150 and an RS/6000 Model F80 to AIX 5L Version 5.1. We then make a system backup by tape and restore it on a partition of an IBM @server pSeries 690. We describe the differences when using a 32-bit und 64-bit POWER3 machine in the following sections.

Migration to AIX 5L Version 5.1

We migrate our test machines from AIX Version 4.3 to AIX 5L Version 5.1 either by CD, as described in this chapter, or by using NIM, as we describe in Chapter 6, “Migration by NIM” on page 105. We update the systems after the migration to AIX 5L Version 5.1 ML 03. To verify that the system is at the highest AIX level, we follow the instructions in Example 5-9 on page 90.

Example 5-8 Verification of maintenance level prerequisites

```
# instfix -i | grep ML
    All filesets for 5.1.0.0_AIX_ML were found.
    All filesets for 5100-01_AIX_ML were found.
    All filesets for 5100-02_AIX_ML were found.
    All filesets for 5100-03_AIX_ML were found.
```

If your maintenance level is not at the required level, get the latest filesets from the following URL by searching for APAR IY32749:

http://techsupport.services.ibm.com/server/aix.fixsearch51

If your hardware is 64-bit, you need to make sure the following filesets are installed on your system:

- devices.chrp_lpar.base.ras
- devices.chrp_lpar.base.rte
- devices.chrp.base.rte
- devices.chrp.base.ServiceRM

Furthermore, you need to run the partition_ready script, which will install all needed devices and kernel support on your AIX 5L Version 5.1 machine, as shown in Example 5-9. This ensures that the operating system image is bootable on a partitioned POWER4 system. When you run the script, you will be prompted to insert the Update CD in your CD-ROM drive. This way, all device drivers are installed, filesets are updated, and support for logical partitions is added, as provided with APAR IY22854.

Example 5-9 Partition_ready script

```
# /usr/sbin/partition_ready
```

Alternatively, manually install APAR IY22854 and the filesets listed above if you have problems running the partition_ready script. Either way, you enable your AIX system to run not just on a server but also on a partitioned server. If you run into problems with the partition_ready script, then the reason will probably be that the CD-ROM is not recognized by the system. Install all filesets manually, as described above, including the latest maintenance level (ML03) and APAR IY22854. This way, you will still be able to restore your system image on a partitioned POWER4 system.

When we tested the above scenario, it worked without any problems with the 64-bit machine, but we run into problems with the 32-bit machine. The restore on the IBM @server pSeries 690 partition stopped on LED 0c43. The reason for this is that we used an AIX 5L Version 5.1 CD dated (5100-01) and then installed ML03 afterwards as required. However, if you use the most recent base AIX CD for the migration to AIX 5L Version 5.1, which already is at the ML03 level (5100-03), it works. Additionally, you need to install the following filesets if you need to use a 64-bit system:

- bos.64bit
- bos.mp64

Your system is now prepared to be restored on a partitioned system. The next step is to take an mksysb backup by tape. We describe this procedure in 5.3.3,

"System backup by tape" on page 91 and the restore procedure in 5.3.4, "Restoring system image on POWER4 partition" on page 94.

5.3.3 System backup by tape

There are three different ways to initiate an mksysb backup to tape of your system data. These include the Web-based System Manager, SMIT, and command line. For a general discussion on backup options and available commands both for your system and your user data, refer to 4.6, "Backup of existing environment" on page 65.

Option 1: Using the Web-based System Manager

Start the process by typing **wsm** in the command line of your AIX 5L Version 5.1 POWER3 system. Select **Backup and Restore**, which will display the Backup Overview and Task menu, as shown below in Figure 5-2 on page 92.

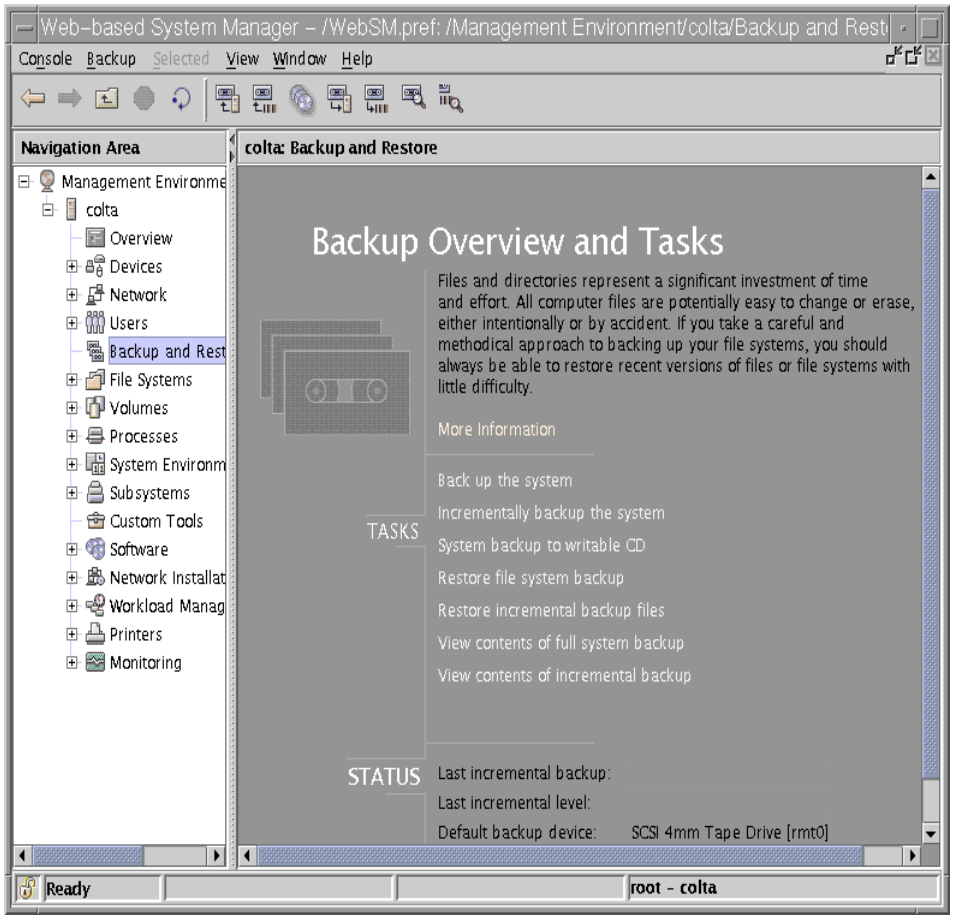

Figure 5-2 Backup Overview and Tasks menu

You then need to select the **Back Up the System** menu. This gets you to the menu shown in Figure 5-3 on page 93. Check the appropriate boxes. You can expand /tmp automatically when the system temporarily needs more space during the backup process. Additionally, you can choose to create map files. However, this is not recommended if you plan to reinstall the backup on a different machine.

Furthermore, you can exclude certain files from the backup, list all files as they are backed up, and generate a new /image.data file. For the backup device, choose the tape device. In our example, this is a 4 mm SCSI tape specified by the device /dev/rmt0.

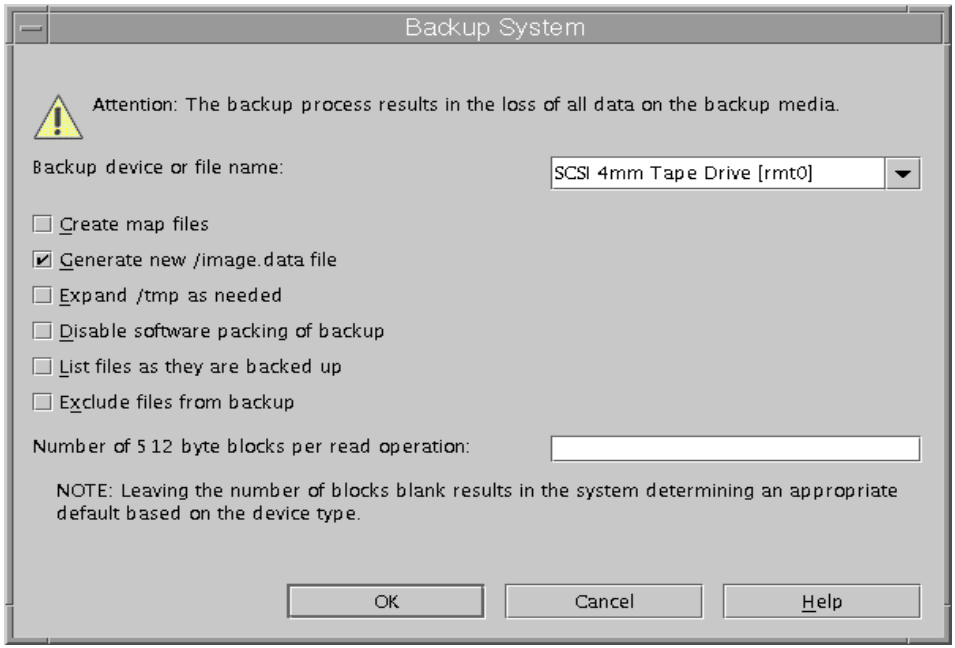

Figure 5-3 Backup onto tape using the Web-based System Manager

Option 2: Using SMIT

Alternatively, you can use SMIT to initiate an mksysb backup onto tape. You can use the **`smitty mksysb`** fast path or just type **`smitty`** and then select **Software Installation and Maintenance** -> **System Backup Manager** -> **Back up the System** to get the screen shown in Example 5-10.

Example 5-10 Backup onto tape using SMIT

```
                          Back Up the System

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

[TOP]                                                   [Entry Fields]
WARNING:  Execution of the mksysb command will
          result in the loss of all material
          previously stored on the selected
          output medium. This command backs
          up only rootvg volume group.

* Backup DEVICE or FILE                                 []  +/
  Create MAP files?                                     no  +
  EXCLUDE files?                                        no  +
```

```
  List files as they are backed up?                    no  +
  Verify readability if tape device?                   no  +
  Generate new /image.data file?                       yes +
  EXPAND /tmp if needed?                               no  +
[MORE...3]

F1=Help             F2=Refresh          F3=Cancel           F4=List
Esc+5=Reset         Esc+6=Command       Esc+7=Edit          Esc+8=Image
Esc+9=Shell         Esc+0=Exit          Enter=Do
```

You can make the same choices described in the Web-based System Manager approach.

Option 3: Using the command line

The third backup option is the command line. This way, we use the **mksysb** command directly, which creates a backup of the operating system. The tape format includes a boot image, a bosinstall image, and an empty table of contents followed by the system backup image.

To initiate the backup, type the following command on the command line:

```
# /usr/bin/mksysb -i /dev/rmt0
```

The -i flag in this example command generates the /image.data file.

5.3.4 Restoring system image on POWER4 partition

In summary, in this first test case, we first migrate a POWER3 32-bit or 64-bit machine from AIX Version 4.3 to AIX 5L Version 5.1. We apply all necessary maintenance levels, filesets, and APARs, as described in “Migration to AIX 5L Version 5.1” on page 89 and run the partition_ready script, which supports the system image on a partitioned system. We back up the system data onto tape using the **mksysb** command. We are now ready to restore the tape with the AIX 5L Version 5.1 system image of our RS/6000 Model 150 and our RS/6000 Model F80 onto a partition of an IBM @server pSeries 690.

For this, we boot the partition at the SMS menu and select the tape drive as the first boot device. This ensures a restoration of the tape on the p690 partition. When this process is finished, you need to configure your network settings and verify that the restore worked without error messages. You are then guaranteed that your system data is successfully transferred from your POWER3 to your POWER4 system. You can then start backing up and restoring your user data. For a general discussion on backing up user data, refer to 4.6, “Backup of existing environment” on page 65.

5.3.5 Case 2: POWER3 to POWER4 AIX 5L Version 5.2 by CD/DVD

In our second test example, we consider migrating an IBM @server pSeries 610, 64-bit hardware, to AIX 5L Version 5.2. We then make a system backup by DVD and restore it on a partition of an IBM @server pSeries 690. We describe the differences from our first test case where we have to install additional filesets and APARs after the migration to make the system image partition ready. The process in this second test case is a lot easier and straightforward because of some features of AIX 5L Version 5.2.

Migration to AIX 5L Version 5.2

The migration of an IBM @server pSeries 610 from AIX Version 4.3 to AIX 5L Version 5.2 works as described in Chapter 5, “Migration by media” on page 79 or as we describe in Chapter 6, “Migration by NIM” on page 105. By default, during a migration installation all device drivers and kernel support is installed on the system. This means that we do not need to prepare our system as described in 5.3.2, “Case 1: POWER3 to POWER4 AIX 5L Version 5.1 by tape” on page 89. Our system image is already partition ready and we can start the backup procedure.

If you decide to have all devices installed at migration time, you need to verify that all necessary device and kernel support is installed on your system to restore it onto another system. We recommend, however, to have the default enabled, which will guarantee that you run into as few problems as possible during your migration and post migrations tasks.

System backup by DVD

There are two ways we can back up our system image on an IBM @server pSeries 610. The first method is by using the **mksysb** command and backup the system data onto tape. We describe this procedure in 5.3.2, “Case 1: POWER3 to POWER4 AIX 5L Version 5.1 by tape” on page 89. Even though we will not go through the whole procedure of how to do it again, we want to confirm that we have still tested this backup option successfully.

Alternatively, you can backup your system data onto CD or DVD. We use DVD in our test case two and will describe the procedure in detail in the following paragraphs. To initiate the backup, you have three methods to choose from, that is, the Web-based System Manager, SMIT, and command line.

For general information on CD-R, DVD-R, or DVD-RAM drives and CD-R, DVD-R, or DVD-RAM creation software, refer to the following README file on your AIX 5L Version 5.2 system at /usr/lpp/bos.sysmgt/mkcd.README.txt.

In this section, we concentrate on backing up an AIX 5L Version 5.2 system by DVD. When creating backups using DVD media, the following two formats are available:

- ISO9660 CD format: Available for DVD-R/DVD-RAM media
- Universal Disk Format (UDF): Available for DVD-RAM media

UDF is a new feature introduced and supported with AIX 5L Version 5.2. UDF uses less disk space for creating the backup because it writes the backup directly to the DVD.

Even though backing up your system by DVD is possible with AIX 5L Version 5.1, there are differences. With AIX 5L Version 5.1, you need to differentiate between backing up a system to reinstall on the same or on another system, as different commands are used: the **mkcd** and the **mkcdgeneric** command. The first one creates an image with only the device drivers needed for this particular system and the second command includes all device drivers. This means with the second option, you can restore the backup on a different server. For additional information on this topic, refer to 4.6, “Backup of existing environment” on page 65.

However, this is not an issue when backing up our AIX 5L Version 5.2 system, because, by default, all device support is installed at migration time. This is what we concentrate on and what we describe in the next sections.

Option 1: Using the Web-based System Manager

The first option to start the backup process on DVD is using the Web-based System Manager. You need to type the **wsm** command and then select **Backup and Restore** in your graphical display. To get to the screen shown in Figure 5-4 on page 97, you need to further select **System Backup by CD/DVD**.

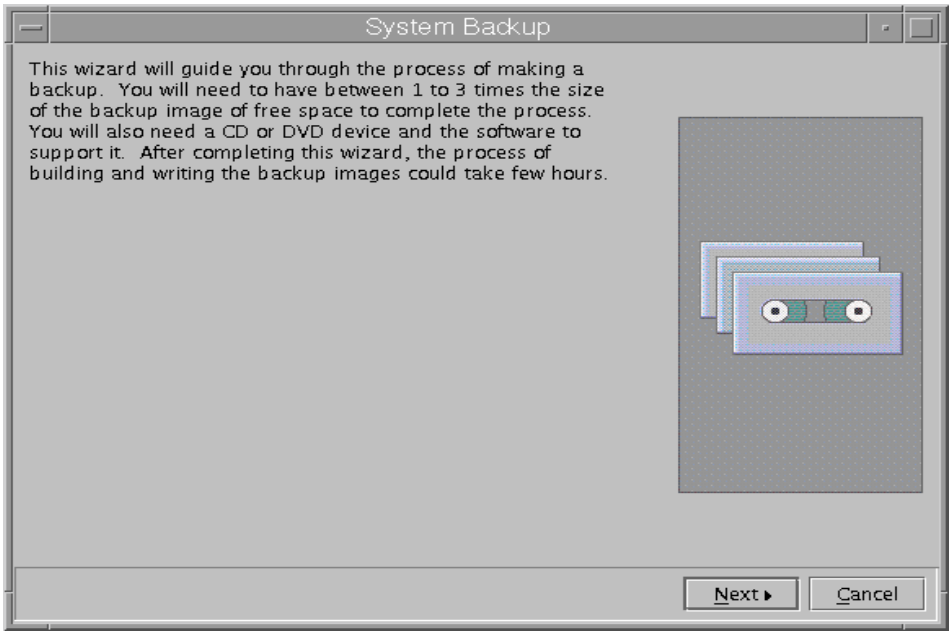

Figure 5-4 Backup by DVD using the Web-based System Manager

You will be guided through the next steps. You need to choose between creating bootable or non-bootable CDs or DVDs and the type of media you will use. Furthermore, you specify whether you want to create a new backup image or whether you want to use an existing backup image. You then need to decide between two different file system formats: ISO9660 (Rockridge) or Universal Disk Format (UDF).

For our test case, we choose the ISO9660 file system format. We also tested the UDF format. When you select UDF file system, you may get the error message shown in Example 5-11.

Example 5-11 UDF file system format error message

```
0512-323 mkcd: The following files are required for the creation of the CD or
DVD image and are not available on the source system:

/usr/lib/drivers/pci/ncr810dd devices.pci.00100100.rte
```

You can solve this problem using the procedure described in Example 5-12 on page 98. You need to delete one line in the specified file. This enables you to use the UDF file system format as a backup option as well.

Example 5-12 Workaround for 0512-323 mkcd error

```
# cd /usr/lpp/bosinst
# vi cdfs.optional.list

Delete the following line and save the file:
/usr/lib/drivers/pci/ncr810dd devices.pci.00100100.rte
```

A summary of our chosen options is shown in Figure 5-5. You can again exclude certain files from the backup, specify temporary working space, and create the image.data and bosinst.data files.

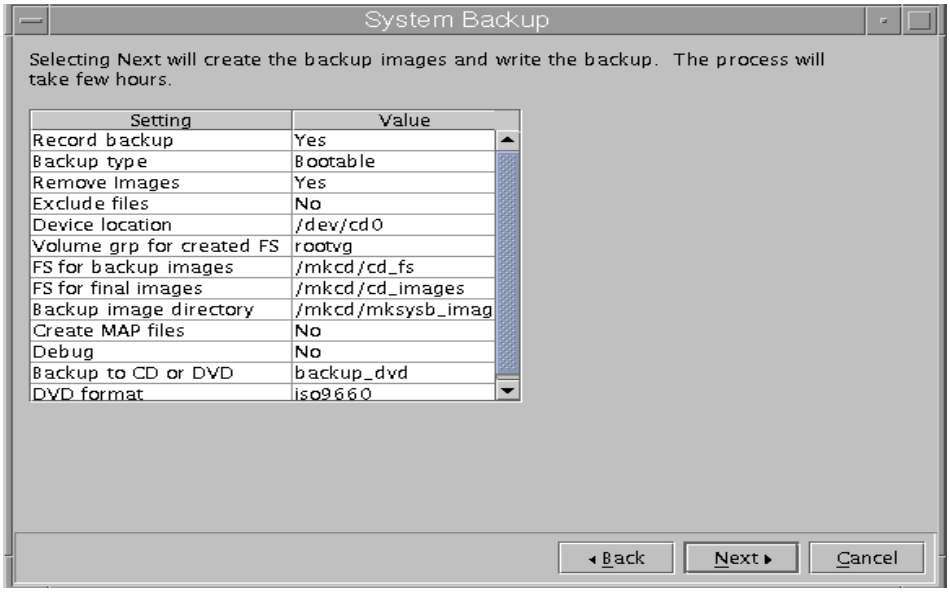

Figure 5-5 Summary of backup by DVD using the Web-based System Manager

Option 2: Using SMIT

Alternatively, you can use SMIT. The fast path for an ISO9660 CD format backup is **`smitty mkcd`** and, for an DVD backup, it is **`smitty mkdvd`**. You get to the same menus by just typing **`smitty`** and then by selecting **System Storage Management** -> **System Backup Manager** -> **Back up the System** -> **Back up this System to DVD**.

When you select the backup to DVD as in our test example, you see the SMIT menu shown in Example 5-13 on page 99 with all the available options.

Example 5-13 Backup by DVD to ISO9660 format using SMIT

```
                    Back Up This System to ISO9660 DVD

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

[TOP]                                                   [Entry Fields]
DVD-R or DVD-RAM Device                                 []  +
mksysb creation options:
Create map files?                                       no  +
Exclude files?                                          no  +

File system to store mksysb image                       []  /
(If blank, the file system
will be created for you.)

File system to store DVD file structure                 []  /
(If blank, the file system
will be created for you.)
[MORE...19]

F1=Help             F2=Refresh          F3=Cancel           F4=List
Esc+5=Reset         Esc+6=Command       Esc+7=Edit          Esc+8=Image
Esc+9=Shell         Esc+0=Exit          Enter=Do
```

Option 3: Using command line

The two commands you need to use when initiating the backup by DVD are shown in Example 5-14. The first line indicates a backup to CD/DVD using the ISO9660 file system format. The second line of this example shows the command to use if you back up the system data to CD/DVD using the UDF file system format.

Example 5-14 Backup by DVD using command line

```
# /usr/sbin/mkcd -L -d /dev/cd0 -V rootvg

# /usr/sbin/mkcd -U -d /dev/cd0 -V rootvg
```

Restoring the system image on POWER4 partition

We are ready to restore the system image taken from an IBM @server pSeries 610 as an example of a 64-bit POWER3 machine with AIX 5L Version 5.2. We boot our IBM @server pSeries 690 partition in SMS mode and configure the CD-ROM as the first boot device. The restore starts when the system boots from DVD.

You need to verify that there are no error messages after the restore process and that you configured your network settings. In the next step, you can transfer your user data onto the partition on your POWER4 machine.

5.3.6 Case 3: Cloning your system

In our third test example, we consider cloning the root volume group of any of our test machines, that is RS/6000 Model 150, RS/6000 Model F80, and the IBM @server pSeries 610 to an alternate disk. We then migrate the cloned disk and make a system backup of the clone by tape or DVD. We then restore it on a partition of an IBM @server pSeries 690 or any other machine. The emphasis in this section is on the alternate disk installation process.

This is an important issue when we discuss one of the new features of AIX 5L: The alternate disk migration procedure in Chapter 6, “Migration by NIM” on page 105. Even though the administrator only issues one command, the **nimadm** command, the procedures behind it are complex and involve a combination of cloning, NIM, and migration. You should therefore understand the issues involved and become familiar with the concept. The alternate disk migration is by far the most complex migration method. It is, however, also the most attractive one with little downtime and a low risk of losing any data or information.

You can either back up your cloned root volume group and restore it onto another machine. Alternatively, you can use the disk and insert it into the target machine. This is a fast and straightforward method that guarantees that all information is transferred to your target machine.

Furthermore, the alternate disk installation allows you to create a copy of your root volume group to one or more free disks. The advantage is that the cloning process takes place while your rootvg is up and running. Another advantage of cloning your rootvg is when you apply fixes. In case these cause problems on your system, you can switch back to the original rootvg, which works without problems. In the meantime, you can solve the problems you encountered with certain fixes.

Prerequisites

Before cloning your system, you need to verify whether or not your system is enabled to be installed on any system. This is one of the choices you have at installation time. If you select to have this option disabled, you will not have all device support installed on your system automatically. This may cause problems when you restore your backup to another hardware. However, you can manually install the missing filesets at by selecting the AllDevicesKernels software bundle to be installed.

To verify that your system is installed with all devices and kernel support, follow Example 5-15.

Example 5-15 Verification of device and kernel support

```
# grep ALL_DEVICES_KERNELS /var/adm/ras/bosinst.data
ALL_DEVICES_KERNELS = yes,no
ALL_DEVICES_KERNELS = yes
```

To check the status of your disks, type the **lspv** command. The output looks similar to the one shown in Example 5-16.

Example 5-16 Availability of disks

```
# lspv
hdisk0          0000317a861060f7                    rootvg
hdisk1          0000317a9c8d0ce7                    None
```

In our test environment, our root volume group is on hdisk0 and we use hdisk1 for the system clone. To initiate the alternate disk installation process, you need to have the bos.alt_disk_install.rte filesets on your AIX machine.

If you involve partitions for your cloning process, you need to apply the following APARs:

- AIX 5L Version 5.2: APAR IY35456 - bos.alt_disk_install.5.2.0.1
- AIX 5L Version 5.1: APAR IY35312 - bos.alt_disk_install.5.1.0.36

As in most of our examples before, there are three methods to clone your system to another disk. These include the Web-based System Manager, SMIT, and command line. However, we will only demonstrate how to clone your root volume group onto another disk using SMIT.

Using SMIT

You can use the **smitty alt_clone** fast path or just type **smitty** and then select t **Software Installation and Maintenance** -> **Alternate Disk Installation** -> **Clone the rootvg to an Alternate Disk** to see the screen in Example 5-17.

Example 5-17 Alternate Disk Installation using SMIT

```
                  Clone the rootvg to an Alternate Disk

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

[TOP]                                                 [Entry Fields]
Target Disk(s) to install                             []  +
```

```
Phase to execute                                        all +
image.data file                                         []  /
Exclude list                                            []  /

Bundle to install                                       []  +
-OR-
Fileset(s) to install                                   []

Fix bundle to install                                   []
-OR-
Fixes to install                                        []

[MORE...18]



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

You are required to specify a target disk for the installation. In our test environment, we use hdisk1. Alternatively, you could install an mksysb on an alternative disk. The fast path is **smitty alt_mksysb**. We will not describe it in detail though because the procedure is very similar to cloning the rootvg to an alternate disk.

You can verify that the alternate disk was created by typing the **lspv** command again, as we show in Example 5-18. Your output should look very similar. In our test environment, hdisk0 is still the active root volume group. The status of hdisk1 has changed. It is now called the altinst_rootvg, which is the cloned rootvg.

Example 5-18 Status of hard disks before reboot

```
# lspv
hdisk0          0009027724fdbd9f            rootvg              active
hdisk1          0009027779fe61c6            altinst_rootvg
```

When you reboot your system to the cloned disk, the **lspv** output looks different. This is shown in Example 5-19 on page 103.

Example 5-19 Status of hard disks after reboot

```
# lspv
hdisk0          0009027724fdbd9f            old_rootvg
hdisk1          0009027779fe61c6            rootvg                  active
```

The cloned rootvg on hdisk1 is now the active rootvg. The status of hdisk0 is now old_rootvg.

5.3.7 An alternative

An alternative way (in addition to the three cases described above) to move from a POWER3 to a POWER4 partition is to migrate your POWER3 AIX Version 4.3 system to AIX 5L, then take this disk and physically insert it into a POWER4 server. This method works and is supported. In this case, you need to execute the **devreset** command to clean out all of the device attributes. Furthermore, you need to make sure that you have all devices and kernel support installed. Either run the partition_ready script, or manually install the missing filesets, as described in “Migration to AIX 5L Version 5.1” on page 89.

6

Migration by NIM

In this chapter, we describe the Network Installation Management (NIM) method of migrating from AIX Version 4.3 to AIX 5L. It is a medialess process that involves many different migration paths, depending on the size and complexity of your environment. Furthermore, we will include detailed information on prerequisites, dependencies, and necessary pre- and post migration tasks. Some of the main concerns and issues of this chapter are as follows:

- Migration of the NIM master
- Migration of NIM clients
- Alternate disk migration
- Minimizing the system downtime
- Availability during migration process
- Troubleshooting

6.1 NIM enhancements

With AIX 5L, there are improved NIM features that will assist you in installing and migrating your AIX Version 4.3 system. The following section describes some of the most important enhancements, which include:

- Easy NIM (EZNIM)
- Alternate disk migration
- Verification of NIM resources
- Improved NIM interfaces
- Service consumability

6.1.1 Easy NIM (EZNIM)

One possibility to set up your NIM environment is to use the simplified SMIT EZNIM menus. It is useful for administrators who have limited or no knowledge about the Network Installation Management. The EZNIM menus include automatic resource naming and offer the user the ability to review the steps that will be initiated before the task is executed. It gives the administrator the ability to perform the most common tasks without much input.

There are only two main categories to choose from: The first choice is setting up the NIM master and the second choice is setting up the NIM client. Furthermore, you only get the most frequently used menu tasks to configure your NIM environment. We will show some detailed screenshots and explain the appropriate menus.

If you familiar with NIM and know how to set up the NIM master and the client machines, you will probably use this approach. You will find some additional SMIT menus and features that we explain in detail in this section, but the main setup has remained unchanged.

To get to the SMIT EZNIM menu, you can either type the **`smitty eznim`** fast path or just type **`smitty`** and then select **Software Installation and Maintenance** -> **EZ NIM (Easy NIM Tool)**.

If your NIM environment has not been set up on your AIX 5L system, the EZNIM menu looks as shown in Example 6-1 on page 107.

Example 6-1 EZNIM main menu

```
                    EZ NIM (Easy NIM Tool)

Move cursor to desired item and press Enter.

    Configure as a NIM Master
    Configure as a NIM Client



F1=Help             F2=Refresh       F3=Cancel        F8=Image
F9=Shell            F0=Exit          Enter=Do
```

This main menu allows you to choose between configuring your machine as a NIM master or as a NIM client. There are several submenus, depending on the choice you make. We explain these in detail in the following sections.

Configure a NIM master

If you want to configure your AIX 5L machine as a NIM Master, select it in the main SMIT EZNIM menu, as shown in Example 6-1, and the options that are available to you are displayed in Example 6-2.

Example 6-2 Configuring a NIM master using EZNIM

```
                    Configure as a NIM Master

Move cursor to desired item and press Enter.

Setup the NIM Master environment
Add fixes to the NIM Master environment
Add client to the NIM environment

Update clients
Backup a client
Reinstall clients
Reset clients

Show the NIM environment
Verify the NIM environment
Remove NIM environment



F1=Help          F2=Refresh       F3=Cancel        F8=Image
F9=Shell         F0=Exit          Enter=Do
```

The first step is to set up the NIM master environment using the first menu in Example 6-2 on page 107. You need to specify the software source needed to configure the NIM resources. Additionally, you need to select the volume group as well as the file system to use for the NIM resources. Executing this step configures your machine as the NIM master and creates the basic resources, such as SPOT, lpp_source, and the bosinst.data file.

You can use the second menu of Example 6-2 on page 107 to update your resources. This includes installing new filesets onto the SPOT resource. Furthermore, you can update all your client machines during this operation at the same time.

Other menus of the main NIM master configuration menu give you the option to add clients to your NIM environment, and to back up, reinstall, or reset client machines. For the backup option, you can initiate a backup of the client from the master and also store the backup on the master machine. Reinstalling a client means you perform an **mksysb** restore on the client, for which you need to specify the location of the backup image.

Using the last three menus gives you the possibility to verify and view the NIM environment that you set up. Alternatively, you could use the **lsnim** command on the NIM master to get the same information. To delete all NIM definitions and resources, you need to select **Remove NIM environment** from the main menu.

Configure a NIM client

The second option you have when using the SMIT EZNIM menu (as shown in Example 6-1 on page 107) is to configure your AIX machine as a NIM client. When you choose this option, you will get the SMIT display shown in Example 6-3.

Example 6-3 Configuring a NIM client using EZNIM

```
                        Configure as a NIM Client

Move cursor to desired item and press Enter.

Add this system to a NIM environment
Update this system
Reinstall this system
Reset this system



F1=Help             F2=Refresh          F3=Cancel        F8=Image
F9=Shell            F0=Exit             Enter=Do
```

You can perform operations on the NIM master like updating, reinstalling, and resetting the client. The last option resets the state of the client and is used in general after a failed NIM operation. It returns the NIM client to the *ready* state. Furthermore, you can add this AIX machine as a client to the existing NIM environment, in this case, it is initiated from the client machine.

6.1.2 Alternate disk migration

If you have tight downtime requirements for your production environment, the alternate disk migration method introduced with AIX 5L will be of great interest to you. To use this enhanced feature, you need a NIM master running AIX 5L Version 5.1 with ML 03 or later with the appropriate NIM resources. Additionally, you need client machines running AIX Version 4.3 or later that have a spare disk or, in the case of a mirrored environment, spare disks. A graphical summary of the setup is shown in Figure 6-1.

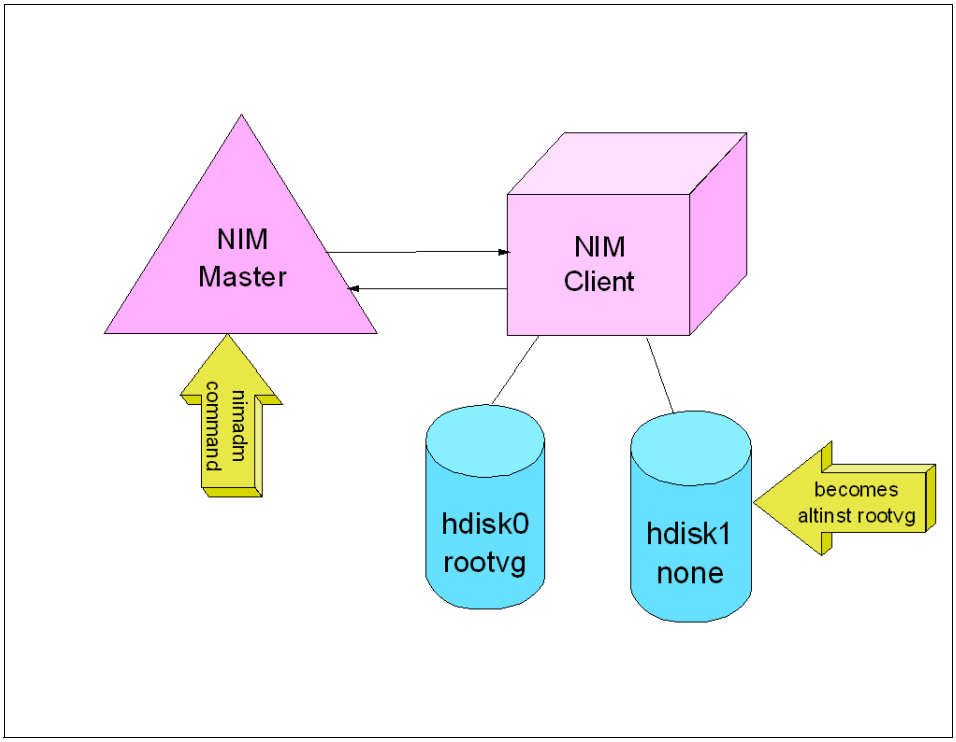

Figure 6-1 Graphical overview of alternate disk migration concept

By issuing the **nimadm** command, either by using SMIT or the command line, you initiate a cloning of your AIX Version 4.3 disk and a simultaneous migration of the cloned disk to AIX 5L. This procedure runs while the root volume group of your

AIX Version 4.3 system is running. When the process is finished, you can reboot the system at a time that least affects your production environment and switch to your cloned and migrated root volume group. The downtime is therefore reduced to the amount of time it takes your AIX server to reboot. It is a method with many more advantages, which will be described in detail, including an example, in 6.1.2, “Alternate disk migration” on page 109.

6.1.3 Improved NIM interfaces

There are two major improvements to the usability of lpp sources with AIX 5L Version 5.2. These include the addition of the **lppmgr** command to manage install images from an lpp source and the addition of an update operation.

The lppmgr command

To help the system administrator manage and maintain their server environments, the **lppmgr** command was introduced with AIX 5L Version 5.2. You can use it to clean up software images either using the SMIT menu or command line. You have the option to remove duplicate updates or base levels. Furthermore, you are able to remove message and locale filesets for languages that you do not require. You can remove superseded filesets as well as non-system images from an lpp source resource.

Using SMIT, you can either use the **smitty nim_lppmgr** fast path or just type **smitty** and then select **Software Installation and Maintenance** -> **Network Installation Management** -> **Perform NIM Administration Tasks** -> **Manage Resources** -> **Perform Operations on Resources** -> **lppsource_aix52** -> **lppmgr** -> **Eliminate unnecessary software images in an lpp_source** to get to the screen shown in Example 6-4.

Example 6-4 Eliminating unnecessary software images in an lpp_source

```
         Eliminate Unnecessary Software Images in an lpp_source

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

                                          [Entry Fields]
TARGET lpp_source                         lppsource_aix52
PREVIEW only?                             yes                    +
REMOVE DUPLICATE software                 yes                    +
REMOVE SUPERSEDED updates                 yes                    +
REMOVE LANGUAGE software                  yes                    +
PRESERVE language                         [en_US]
REMOVE NON-SIMAGES software               no                     +
SAVE removed files                        no                     +
DIRECTORY for storing saved files         []
```

```
EXTEND filesystems if space needed?            yes                    +


F1=Help             F2=Refresh       F3=Cancel        F4=List
F5=Reset            F6=Command       F7=Edit          F8=Image
F9=Shell            F0=Exit          Enter=Do
```

You can also use the command line. The syntax of the **lppmgr** command is:

```
# nim -o lppmgr -a lppmgr_flags=<flags> <lpp_source_object>
```

The update operation

The second enhancement in NIM interfaces is the addition of an update operation. With AIX Version 4.3 and AIX 5L Version 5.1, you would copy packages into an lpp directory or remove them and run the **nim -o check** command to get the lpp_source attributes updated. This is not necessary with AIX 5L Version 5.2. The improved update operation performs all necessary checks automatically.

As with the first example of improved NIM interfaces, you can use the SMIT menus or command line to update the lpp_source attributes. When using SMIT, run the **smitty nim_update** fast path or just type **smitty** and then select **Software Installation and Maintenance** -> **Network Installation Management** -> **Perform NIM Administration Tasks** -> **Manage Resources** -> **Perform Operations on Resources** -> **lppsource_aix52** -> **Update- add or remove software to or from an lppsource**.

An example of how to add software to an existing lpp source is shown in Example 6-5. You need to specify the target lpp source, the software source, and the software packages or install bundles to add.

Example 6-5 Add software to an existing lpp_source

```
                    Add Software to an lpp_source

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

                                                      [Entry Fields]
TARGET lpp_source                                     lppsource_aix52
SOURCE of Software to Add                             /dev/cd0
SOFTWARE Packages to Add                              [all]
INSTALLP BUNDLE containing packages to add            []
gencopy Flags
DIRECTORY for temporary storage during copying        [/tmp]
EXTEND filesystems if space needed?                   yes
```

```
F1=Help              F2=Refresh       F3=Cancel        F4=List
F5=Reset             F6=Command       F7=Edit          F8=Image
F9=Shell             F0=Exit          Enter=Do
```

The software to add can only be processes from a single-volume CD. If you want to add filesets that are on different CDs, you need to copy them first into a file and choose this field as the software source location. To remove software from an existing lpp_source, use the same SMIT menus described above, but choose **Remove Software** rather than **Add Software**.

Using the command line requires the syntax shown in Example 6-6. In the first case, we want to add the bos.perf.tune fileset from CD to the lpp source called lppsource_aix52. In the second case, we want to remove this fileset again from the same lpp source.

Example 6-6 Update operation using the command line

```
# nim - o update -a packages="bos.perf.tune" -a source=/dev/cd0 lppsource_aix52

# nim -o update -a packages="bos.perf.tune" -a rm_images=yes lppsource_aix52
```

6.1.4 Verification of NIM resources

With AIX 5L Version 5.2, you have the ability to verify the validity of your NIM resources. This includes the verification of your SPOT, lpp_source, and mksysb resources. This way, you are more flexible when working with resources that are used during a bos_inst operation.

To get to the SMIT **Verifying resources** screen, type **`smitty`** and then select **Software Installation and Maintenance** -> **Network Installation Management** -> **Perform NIM Administration Tasks** -> **Manage Resources** -> **Verify Resources**.

At this stage, you can choose between an existing SPOT, lpp_source, and mksysb resource. When choosing the SPOT, both the **`lppchk -v`** and **`lppchk -c`** commands will be used. This guarantees that the /usr and the /(root) part of the SPOT are consistent with each other. It also verifies that all prerequisites of the installed software are met. The SMIT screen for verifying an lpp_source is similar to Example 6-7 on page 113.

Example 6-7 Verifying an lpp_source

```
                    Verify an lpp_source resource

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

                                                      [Entry Fields]
Resource Name                                         lppsource_aix52
Perform SIMAGES recovery?                             no                +
Source of install images to be used                   []                +



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

When you choose to verify your lpp_source resource, the system performs a check to determine whether or not the lpp_source contains all the filesets required to meet the images requirements.

6.1.5 Service consumability

Another enhancement of AIX 5L Version 5.2 NIM is called service consumability, which provides customers with an easier mechanism to update their systems with fixes and maintenance levels. You can use this new feature using SMIT or the command line. There are four comparison reports available, which include the following:

- Compare installed software to an image directory (fix repository)
- Compare installed software to a list of updates available from the support Web site
- Compare fix repositories to a list of updates available from the support Web site
- Compare installed software on a base system to another system

The last comparison report is available by the command line only. The SMIT fast path to the main comparison menu is **`smit service_software`**. You can find a detailed description, including screen shots, in the *AIX 5L Differences Guide Version 5.2 Edition*, SG24-5765.

6.1.6 Additional enhancements

There are some further enhancements with Network Installation Management in AIX 5L that we will discuss briefly in this section. You will find additional information in 6.3, “NIM master migration” on page 117 and in 6.4, “NIM client migration” on page 137.

Simultaneous resource creation

With AIX 5L Version 5.1 and earlier, you could only create one resource at a time. The reason is that you have a lock on your server for the duration of the resource creation. With AIX 5L Version 5.2, this limitation has been lifted. You can now create multiple lpp_sources and mksysb resources to a NIM server simultaneously.

The lock is not on the server anymore, but only on one file system, which allows multiple resources to be created that use different file systems. However, if your file system is large enough, you can still create multiple resources simultaneously under the same file system. To do this task, use the force option, which prevents the locking mechanism.

One limitation is, however, that you cannot create simultaneously multiple SPOTs. Nevertheless, you can create lpp_sources, a SPOT, and mksysb resources at the same time.

The mksysb restore

With AIX 5L Version 5.2, all device support is installed by default. This means that for an mksysb restore the lpp_source is not needed anymore. This presumes that you did not change the default installation value, so you have all device support added to your mksysb image.

Network install interface

Another new feature of AIX 5L Version 5.2 is that you can add speed and duplex values for the **`bootlist`** command. These settings will only apply to clients on an Ethernet network. You obtain the SMIT screen shown in Example 6-8 on page 115 by typing **`smitty`** and then selecting **Software Installation and Maintenance** -> **Network Installation Management** -> **Perform NIM Administration Tasks** -> **Manage Machines** -> **Define a Machine** or **Change/Show Characteristics of a Machine**.

Example 6-8 Speed and duplex values of the network interface

```
                                Define a Machine

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

                                                      [Entry Fields]
NIM Machine Name                                      [srvr80j]
Machine Type                                          [standalone]       +
Hardware Platform Type                                [chrp]             +
Kernel to use for Network Boot                        [up]               +
Primary Network Install Interface
Cable Type                                             bnc               +
Network Speed Setting                                 []                 +
Network Duplex Setting                                []                 +
NIM Network                                           [ent-Network1]
Network Type                                          ent
Ethernet Type                                         Standard           +
Subnetmask                                            []
Default Gateway Used by Machine                       []
Default Gateway Used by Master                        []
Host Name                                             srvr80j
Network Adapter Hardware Address                      [0]
Network Adapter Logical Device Name                   []
IPL ROM Emulation Device                              []                  +/
CPU Id                                                []
Machine Group                                         []                  +
Comments                                              []



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

The possible options for the speed settings are 10, 100, and 1000. For the duplex values, you can choose between auto, full, and half.

6.2 NIM environment

Before you start the migration of your server environment, one of the first steps for you to check is whether or not your hardware is still supported with AIX 5L. Customers with Micro Channel Architecture (MCA) and PowerPC (PReP) architecture machines especially need to carefully plan how to port their applications onto newer machines. The reason is that with AIX 5L Version 5.2,

the only supported hardware is CHRP machines. You can find more information on supported hardware components in Chapter 4, “Planning and preparation” on page 37. We also include a detailed list of withdrawn hardware with AIX 5L Version 5.1 and Version 5.2 in Appendix A, “Hardware and fileset support with AIX 5L” on page 197.

6.2.1 Defining the NIM environment

The positive side of moving to newer and more powerful machines is that you can use features like logical and dynamic partitioning as well as enhanced Workload Manager features. These guarantee you a better usage of the resources in your system and can also increase your performance. Most machines supported with AIX Version 4.3.3 are still supported with AIX 5L Version 5.1. This means that even if you have hardware that is not supported with AIX 5L Version 5.2, you can still migrate your system, but only to AIX 5L Version 5.1. However, it does offer an attractive alternative and ensures that you can still keep your older machine in your environment as long as AIX 5L Version 5.1 is supported.

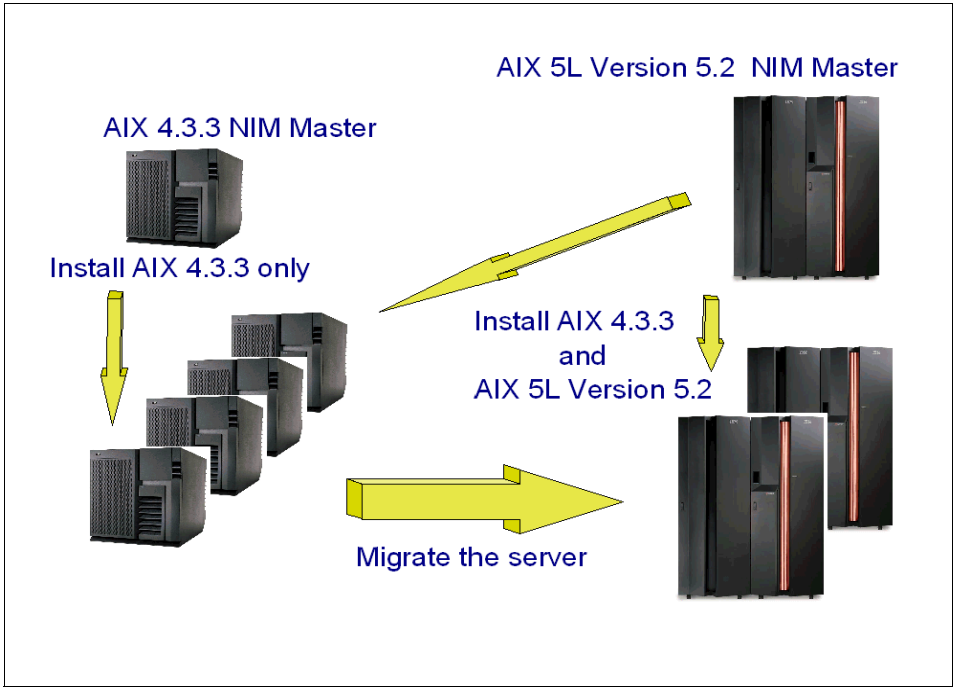

Figure 6-2 The NIM migration environment

In the following sections, we demonstrate with several examples different ways of migrating your environment via the network. Our example setup might be

similar to your environment, which will help you to plan and prepare your AIX machines for the migration process. However, in some aspects your environment might be more complex or different. The important issue is to give you insight into the different NIM migration methods. We will include a detailed explanation of how each one of them works, when it should be used, and what restrictions, dependencies and special tasks you have to do before or during the migration process. This should help you to run into as few problems as possible using the NIM migration method in your production environment.

The order in which you migrate your machines matters. This means that you should first consider the NIM master and then the NIM clients. The NIM master needs to run on the highest operating system level and hence needs to be migrated first. That is also where we will start our discussion. In 6.3, “NIM master migration” on page 117, we discuss NIM master related issues followed by a discussion and some example setups of the NIM client issues in 6.4, “NIM client migration” on page 137. These include three ways to get your NIM client to an operating system version of AIX 5L Version 5.2: network new and complete overwrite, network migration, and alternate disk migration.

6.2.2 Operating system requirements

Before we start with the NIM master migration process from AIX Version 4.3 to AIX 5L Version 5.2, it is important to specify the requirements regarding the AIX level of your AIX Version 4.3 system. The latest maintenance level of AIX Version 4.3 is ML 10. However, not every customer has each machine at this level. Many of you will have a variety of maintenance levels of AIX Version 4.3 in your environment. Some may include additional fixes, PTFs (Program Temporary Fixes) or APARs (Authorized Program Analysis Reports), which may be missing on other machines.

To always have an updated system and to be able to use the latest improvements in code, we advise you to have your system at AIX Version 4.3 ML 10. For the migration, however, it is not a prerequisite. The migration process is designed in a way that you can start your migration at any level of AIX Version 4.3. We have not tested every single maintenance level, but we provide you with details of which level of AIX Version 4.3 we have tested in each test scenario.

6.3 NIM master migration

There are several options to migrate your NIM master from AIX Version 4.3 to AIX 5L. One option is to migrate your original NIM master directly from AIX Version 4.3 to AIX 5L using media, as described in Chapter 5, “Migration by media” on page 79. Using media is the only migration method for a NIM master.

The drawback of this approach is, however, that during the migration process the machine will be unavailable as NIM master. This includes unavailability for any disaster recovery, which may be an important issue in your environment.

Another option is to newly install a different machine with AIX 5L Version 5.2 and then to set up this machine as the new NIM master. The advantage is that the running AIX Version 4.3 NIM master keeps operating while the new NIM master is being installed. You can install the AIX Version 4.3.3 lpp_source on the AIX 5L Version 5.2 NIM master and let the master create an AIX Version 4.3.3 SPOT. This ensures that the new AIX 5L Version 5.2 NIM master is able to install the backups of any of your AIX Version 4.3 systems. This works in a similar way with AIX 5L Version 5.1. This means that the new NIM master can act as a NIM server for AIX Version 4.3 and both versions of AIX 5L.

This method is in particular attractive if you have an existing NIM master on hardware that is not supported with AIX 5L Version 5.2. This way, you can keep this server as NIM master for AIX Version 4.3 and use another machine as the NIM master for AIX 5L Version 5.1 or 5.2.

Regardless of the method you will use to migrate your NIM master, there are some requirements for the NIM master. It needs to be a stand-alone machine and must be able to communicate with all the machines on the network. Furthermore, it needs to be able to remotely run commands to install the other machines in the environment. The /etc/hosts file must contain entries of the client machines on the environment.

6.3.1 Environment setup

Before we describe the migration of the NIM master from AIX Version 4.3 to AIX 5L Version 5.2, it is important that we define our test environment and give you details of our hardware and software configuration. For our test scenario, we use an RS/6000 Model F50 CHRP machine as the NIM Master. As described in Chapter 4, “Planning and preparation” on page 37, we first update all the system microcode to the latest level. This is recommended before starting the migration procedure of any machine in your server environment. The network type we use for this example migration is a standard Ethernet network.

Our NIM Master initially has been configured to support an AIX Version 4.3.3 environment. The resources are listed in Example 6-9 on page 119 and include the lpp_source433 and a SPOT433. The sample output is obtained by typing the `lsnim` command.

Example 6-9 lsnim output of AIX Version 4.3 NIM master

```
master                  machines                master
boot                    resources               boot
nim_script              resources               nim_script
ethernet_1              networks                ent
lpp_source433           resources               lpp_source
App-Dev                 resources               installp_bundle
Client                  resources               installp_bundle
Pers-Prod               resources               installp_bundle
Server                  resources               installp_bundle
spot433                 resources               spot
srvr80e                 machines                standalone
```

Saving the NIM database files

To save the NIM database before starting the migration from AIX Version 4.3.3 to AIX 5L Version 5.2, you can use the NIM database backup function. This backup will copy the NIM database files and /etc/niminfo into the /etc/NIM.level on the NIM master. The content of this file will be used to compare the NIM version during the NIM database restore. The NIM database restore can only be done on a NIM master with the same NIM fileset versions.

Option 1: Using the Web-based System Manager

To back up the NIM database using the Web-based System Manager, refer to Figure 6-3 on page 120. Begin by typing the **wsm** command. Then select the following:

1. Select **NIM**.
2. Select **NIM** -> **Back Up Database** in the menu.

Additionally, you need to specify the device or file where to create the backup and start the process.

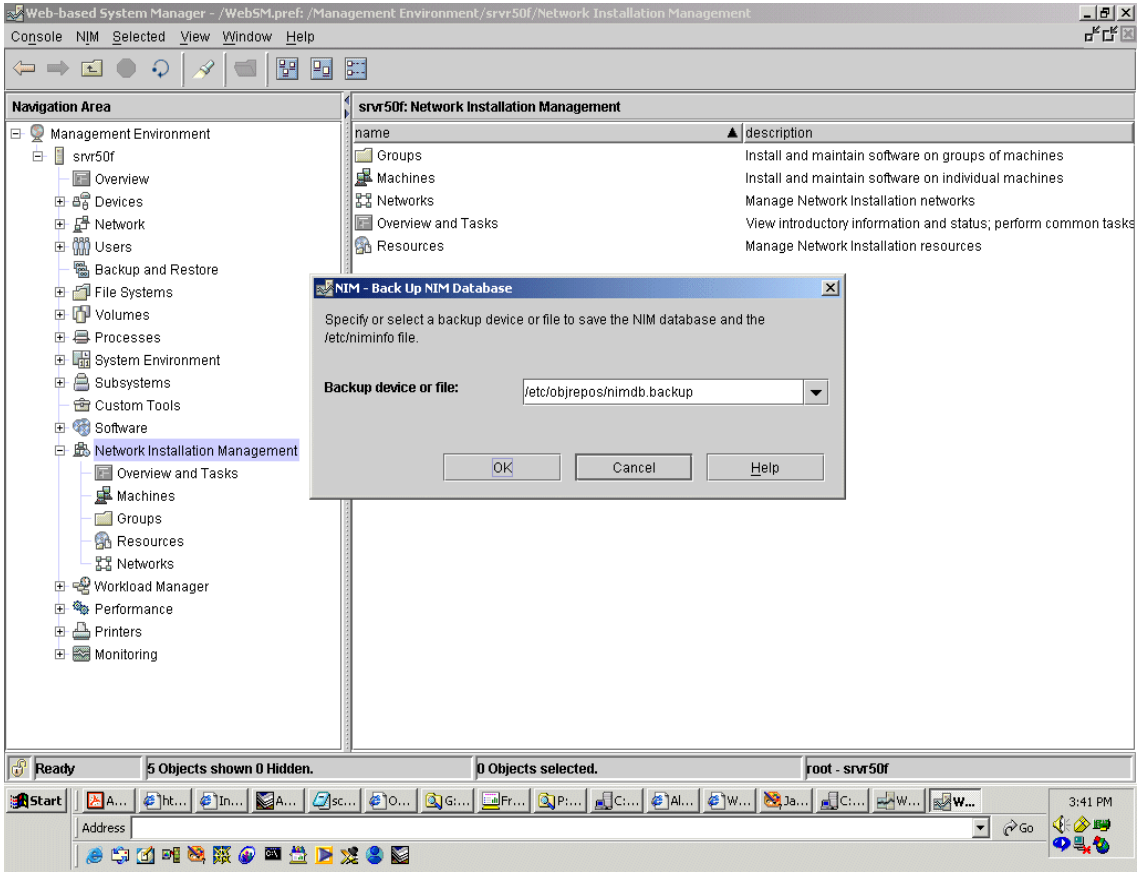

Figure 6-3 Creating a NIM database backup using Web-based System Manager

Option 2: Using SMIT

Alternatively, you can use SMIT to back up the NIM database. Either use the **`smitty nim_backup_db`** fast path to get to the screen shown in Example 6-10 on page 121 or just type **`smitty`** and then select **Software Installation and Maintenance** -> **Network Installation Management** -> **Perform NIM Administration Tasks** -> **Backup/Restore the NIM Database** -> **Backup the NIM Database**.

The default file name to store the backup with is /etc/objrepos/nim.backup.

Example 6-10 Creating the NIM database backup using SMIT

```
                        Backup the NIM Database

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

                                                    [Entry Fields]
* Filename/Device for the Backup                    [/etc/objrepos/nimdb.backup]



F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F0=Exit             Enter=Do
```

Option 3: Using the command line

The third option to back up the NIM database is the command line. You need to save the files listed below. One possibility is to use the **tar -cvf <device_name>** command:

```
/etc/niminfo \
/etc/objrepos/nim_attr \
/etc/objrepos/nim_attr.vc \
/etc/objrepos/nim_object \
/etc/objrepos/nim_object.vc
```

Important: The NIM database backup should only be restored onto a machine that is at the same or higher NIM fileset level than the level of the NIM database backup filesets.

To check the NIM master fileset level, enter to the **lslpp** command, as shown in Example 6-11, which also shows the test machine results.

Example 6-11 Level of NIM master filesets

```
# lslpp -l | grep bos.sysmgt.nim.master

bos.sysmgt.nim.master      4.3.3.76  APPLIED    Network Install Manager
bos.sysmgt.nim.master_gui  4.3.3.0   COMMITTED  Network Install Manager - GUI
```

In our case, we can only restore the NIM database on an AIX machine, which runs on AIX Version 4.3 ML 08 or higher.

6.3.2 Migrating the NIM master from AIX Version 4.3 to AIX 5L

Even though there are several migration methods that you can use to migrate your server environment, for the NIM master, the only possible way is to do a migration by media. This is described in detail, including all available choices at migration installation time, in Chapter 5, “Migration by media” on page 79.

However, we will discuss the issues involved after the migration of the NIM master in the next sections of this chapter. We will concentrate on how to set up the migrated machine as a NIM master in AIX 5L Version 5.2. This is important if you plan to use this AIX server to migrate all other machines in your environment.

Another option is to migrate the NIM master configuration files from one machine running AIX Version 4.3 to another machine where AIX 5L is already installed. This would be the case when you have a POWER3 NIM master that you want to move to a POWER4 partition. There are several ways to achieve this task, which we explain, together with various test examples, in 6.3.4, “Moving the NIM master from POWER3 to POWER4” on page 127.

6.3.3 Configuring the NIM master after AIX migration

The starting point in this section is that we have migrated our AIX Version 4.3 NIM master to AIX 5L Version 5.2 and that we have completed all post migration tasks specific to the NIM master, as described in 7.2, “NIM master specific tasks” on page 173. The AIX migration on a NIM master will not directly affect the NIM environment. However, there are some additional tasks to do before configuring this server as AIX 5L Version 5.2 NIM master for the environment.

After you migrate your NIM master to AIX 5L Version 5.2, the new AIX 5L Version 5.2 NIM resources must be created. There is no need to unconfigure the older AIX resources. The NIM master supports resources of AIX Version 4.2 and later. This means that you can use this one machine to serve as the NIM master for various AIX levels, starting with AIX Version 4.2 up to AIX 5L Version 5.2.

Prerequisites

The prerequisites depend on which NIM resources you want to create and whether it will be stored on the NIM master machine. We describe how to configure a AIX 5L Version 5.2 NIM master to support a base operating system installation (bos_inst).

To create a default AIX 5L Version 5.2 lpp_source, you need about 600 MB of disk space. When using the command line to initiate the creation of an lpp_source, you must ensure that you have enough space in your file system. The reason is that you cannot choose to automatically have the file system expanded if more space is required using this method. It is therefore

recommended that you store the image file in a file system when using media as the source. This ensures that the file system gets expanded automatically if the creation process requires additional space.

Creating basic installation resources

After you migrate your NIM master to AIX 5L Version 5.2, the new AIX lpp_source and SPOT must be created to perform future bos_inst installation. These resources are the basic resources needed to install or migrate a NIM client machine. For further and general information about NIM resources, refer to the *AIX 5L Version 5.2 Installation Guide and Reference*, SC23-4389, which can be found at:

http://publib16.boulder.ibm.com/cgi-bin/ds_form?lang=en_US&viewset=AIX/

lpp_source creation

The process of creating an lpp_source basically involves copying the installation images from CD-ROM to your hard disk. One possibility is to copy the content of the six AIX 5L Version 5.2 CDs to one of your disks using the **bffcreate** command and then to create the lpp_source from this directory.

If you do not have sufficient space on the NIM master machine, you can also use the CD image mounted at /dev/cd0 for the creation of the lpp_source. Whichever way you choose, the lpp_source must contain all the required filesets in order to successfully perform a migration installation.

In our test environment, we use the mounted CD to create the lpp_source. You can initiate the process using the Web-based System Manager, SMIT, or command line.

Option 1: Using the Web-based System Manager

To create an lpp_source using the Web-based System Manager, type the **wsm** command and then follow these steps:

1. Select the **Network Installation Management** icon.
2. Select **Resources** under Network Installation Management.
3. Select **Resources** -> **New** -> **Resource** in the menu; you will see the **Add New Resource** dialog.
4. Select **Advanced** -> **lpp_source**.

Option 2: Using SMIT

The fast path for creating an lpp_source using SMIT is **`smitty nim_mkres`**. The detailed path shown below goes to the same screen:

1. Type the **`smitty`** command.
2. Select **Software Installation and Maintenance** -> **Network Installation Management** -> **Perform NIM Administrating Tasks** -> **Manage Resources** -> **Define a Resource** -> **lpp_source**.

As shown in Example 6-12, you need to first choose lpp_source as the resource type to be created. Then fill in the resource name, resource type, server name, location of resource, and source of install images, similar to our test example.

Example 6-12 Creation an lpp_source using SMIT

```
                          Define a Resource

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

                                                      [Entry Fields]
* Resource Name                                       [lpp_source52]
* Resource Type                                       lpp_source
* Server of Resource                                  [master]
* Location of Resource                                [/export/lpp_source]
  Architecture of Resource                            []
  Source of Install Images                            [cd0]
  Names of Option Packages                            []
  Comments                                            []


F1=Help                 F2=Refresh        F3=Cancel         F4=List
F5=Reset                F6=Command        F7=Edit           F8=Image
F9=Shell                F0=Exit           Enter=Do
```

In our example, only the first AIX 5L Version 5.2 base CD will be used to create the lpp_source. This will be stored in /export/lpp_source under the file name lpp_source52.

Option 3: Using command line

Similarly, we can create an lpp_source using the command line. The syntax of the **nim -o define** command is displayed below:

```
-t lpp_source \
-a location=<directory that will contained the installation images> \
-a server=<master or machine client name> \
-a source=<device that contain the installation media>
```

Important: If you are using the command line to create the lpp_source, remember that it does not automatically expand the file system when it needs additional space.

SPOT creation

The SPOT is used to support machines that boot over the network. There are two types of SPOT you can create: The /usr SPOT and the non-/usr SPOT. When you define a SPOT using the /usr, this resource will contain all the software installed in the NIM master machine. This spot is called /usr spot, and does not use as much space on your machine. However, you cannot update it with optional software once it is created.

The other SPOT type is non-/usr spot and is located in a different file system from /usr. It consumes more disk space, but you can install and update software to support your environment. After having created the AIX 5L Version 5.2 lpp_source, you can use it to create a non-/usr spot. This will contain images to support a new AIX 5L Version 5.2 installation or migration.

As seen with many examples before, there are three ways to initiate the creation of the SPOT on your AIX master machine. These are the Web-based System Manager, SMIT, and command line. They are outlined in the next sections.

Option 1: Using the Web-based System Manager

The menu path with the Web-based System Manager is the same as the creation of the lpp-source, but choose the creation of SPOT this time and fill in the required boxes.

Option 2: Using SMIT

Similarly for SMIT, if you want to create the SPOT, follow the same menus or fast paths as for the creation of the lpp_source. The menu is displayed in Example 6-13 on page 126.

Example 6-13 Creation of a SPOT resource using SMIT

```
                                Define a Resource

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

                                                      [Entry Fields]
* Resource Name                                       [spot52]
* Resource Type                                       spot
* Server of Resource                                  [master]              +
* Source of Install Images                            [lpp_source52]        +
* Location of Resource                                [/export/spot]        /
  Expand file systems if space needed?                yes                   +
  Comments                                            []

  installp Flags
  COMMIT software updates?                            no                    +
  SAVE replaced files?                                yes                   +
  AUTOMATICALLY install requisite software?           yes                   +
  OVERWRITE same or newer versions?                   no                    +
  VERIFY install and check file sizes?                no                    +


F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F0=Exit             Enter=Do
```

Option 3: Using the command line

To perform a SPOT creation from the command line, use the command shown below, including the appropriate flag options:

```
nim -o define -t spot
-a location=<directory that will contain the image, or /usr to an /usr-spot>
-a server=<master or the machine name if it is a client>
-a source=<lpp_source or device that contains the installation images>
-a auto_expand=<to automatically expand the file system>
spot_name
```

Verifying resource creation

We are now at a stage where we have migrated our NIM master from AIX Version 4.3 to AIX 5L Version 5.2. We completed all post migration tasks and configured the AIX 5L Version 5.2 lpp_source and SPOT on our NIM master machine. To verify that the creation is successfully finished, we need to check the usability and the status of resources using the **`nim -o check`** command. It

verifies that all necessary filesets are installed in a lpp_source and updates the table of contents in the file location.

We show in Example 6-14 how you can verify the creation of your SPOT and lpp_source. With the command in the first line, you can list the file sets in the lpp_source. The second and third line verify that the creation of the lpp_source and SPOT was successful.

Example 6-14 Verifying resource creation

```
# nim -o showres <lpp_source_name>

# nim -o lppchk <lpp_source_name>

# nim -o lppchk <spot_name>
```

6.3.4 Moving the NIM master from POWER3 to POWER4

In this section, we discuss the issues involved when moving your NIM master from a POWER3 to a POWER4 hardware and simultaneously migrate it to AIX 5L. We will concentrate on the specific NIM master tasks to do. This includes three different ways to migrate and move from POWER3 to POWER4 architecture as shown:

- Case 1:
  - Back up NIM database of AIX Version 4.3 system
  - Restore NIM database on newly installed AIX 5L POWER4 partition
- Case 2:
  - Migration of a POWER3 NIM master to AIX 5L
  - Back up NIM database of AIX 5L system
  - Restore NIM database on newly installed AIX 5L POWER4 partition
- Case 3:
  - Migration of a POWER3 NIM master to AIX 5L
  - Backup system data by tape or CD/DVD
  - Restore system data on POWER4 partition

You can find the general discussion on this topic in 5.3, “Migration from POWER3 to a POWER4 partition” on page 87. We will concentrate only on the special case, that is, if your machine is the NIM master of your server environment.

Case 1: AIX Version 4.3 NIM database restore

The first option you have when moving the NIM master currently running on a POWER3 machine to a POWER4 partition with AIX 5L is by backing up the NIM database of your AIX Version 4.3 machine only. You then initiate a new installation of AIX 5L on your partition and restore the NIM master configuration files of your AIX Version 4.3 system. One thing to remember is that if your NIM

master uses a different host name and network settings, you need to reconfigure it after the restore process.

If your NIM resources are stored on external disks, you can transfer the NIM database by sharing the disks between machines. You need to export the volume group that contains the NIM resources on the AIX Version 4.3 NIM master and import this volume group on the AIX 5L Version 5.2 NIM master. The commands you use for this operation are **exportvg** and **importvg** and are described in detail in the AIX 5L Version 5.2 Command Reference, Volume 1 to Volume 6, found at:

http://publib16.boulder.ibm.com/cgi-bin/ds_rslt#1

Our test setup is shown in Figure 6-4 on page 129. Our POWER3 NIM master machine running AIX Version 4.3 is an RS/6000 Model F80. We want to transfer the NIM master configuration file to a new installed partition of an IBM @server pSeries 690 running AIX 5L Version 5.2.

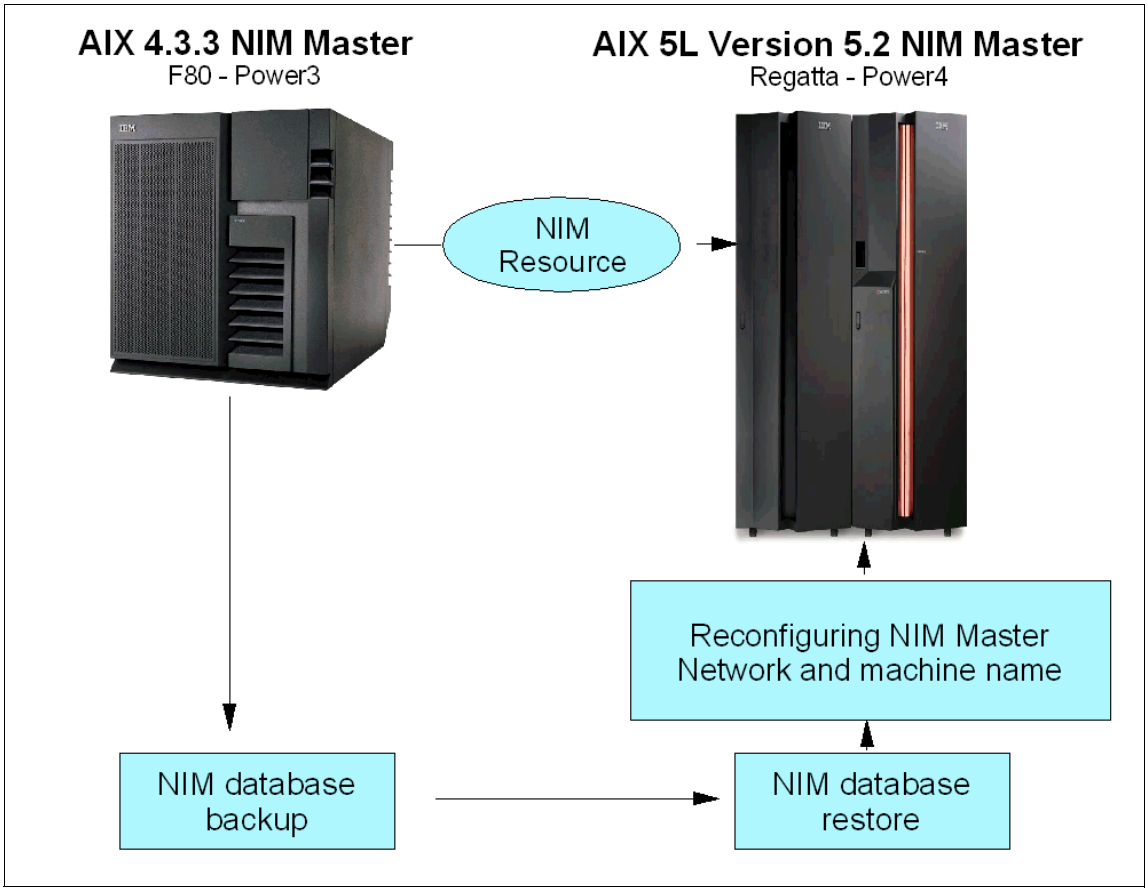

Figure 6-4 Case 1: AIX Version 4.3 NIM database restore

For this setup, we need to first install a partition on an IBM @server pSeries 690 with AIX 5L Version 5.2 using media. Additionally, the POWER4 partition must be able to communicate with all other machines in the NIM environment. The NIM master filesets need to be installed as described in Example 6-11 on page 121 and are listed below:

- bos.sysmgt.nim.master
- bos.sysmgt.nim.spot

The next step is to back up the NIM database on your AIX Version 4.3 machine. This is described in “Saving the NIM database files” on page 119. You can either use the Web-based system manager, SMIT, or command line to initiate the backup process.

When you want to restore the NIM database files, the target machine must be at the same or higher level of AIX as the machine where the NIM database was created. This prerequisite is fulfilled in our test case. We created a NIM database file on the original NIM master running AIX Version 4.3 and want to restore this database file on a partition running AIX 5L Version 5.2 on an IBM @server pSeries 690. The /etc/NIM.level file contains this fileset information.

There are three different ways you can initiate the restore process of the NIM database onto your AIX 5L Version 5.2 partition. We outline all three methods: the Web-based System Manager, SMIT, and command line.

Option 1: Using the Web-based System Manager

To restore the NIM database using the Web-based System Manager, follow the outlined steps to get to the menu shown in Figure 6-5 on page 131. First, type the **wsm** command and select the following:

1. Select **Network Installation Management**.
2. Select **NIM** -> **Restore Database** in the menu.

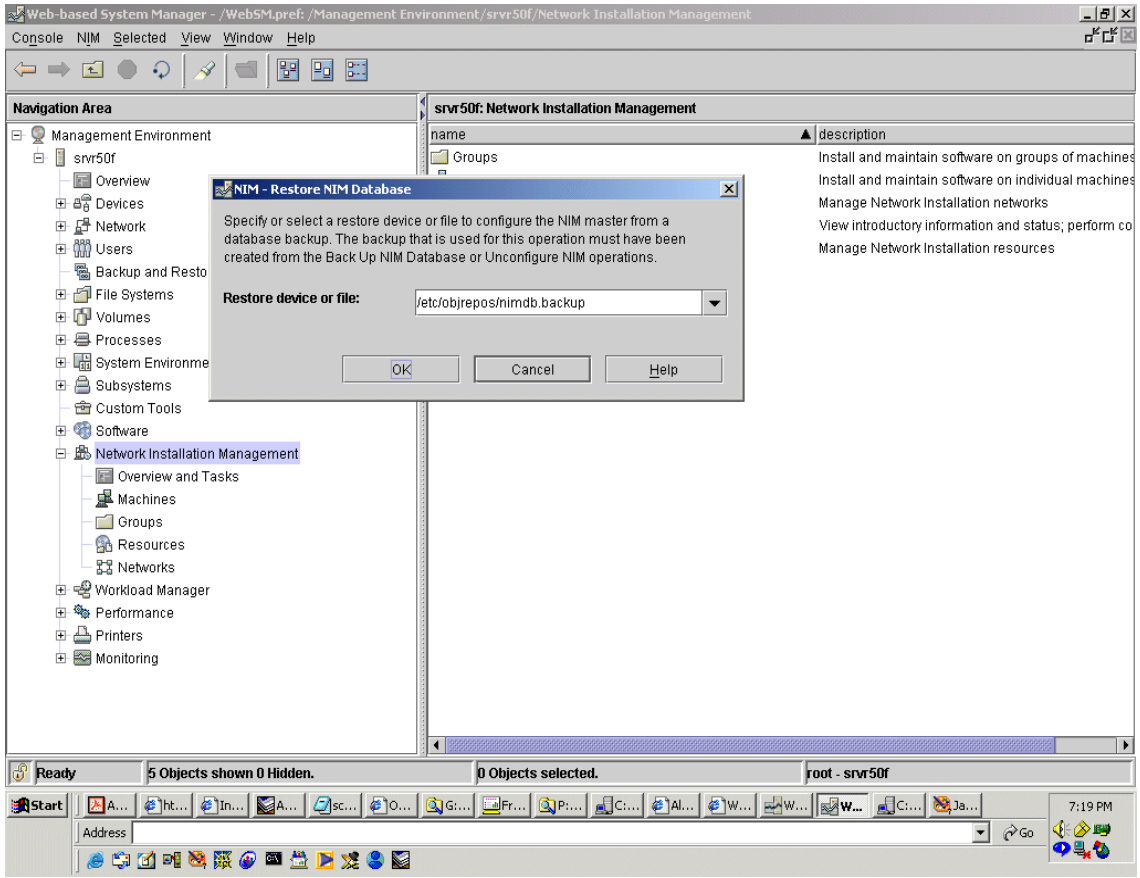

Figure 6-5 Restore of NIM database using the Web-based System Manager

Option 2: Using SMIT

To restore a NIM Database using SMIT, you need to fill in the menu shown in Example 6-15 on page 132. Either type the **`smitty nim_restore_db`** fast path or just type **`smitty`** and then select **Software Installation and Maintenance** -> **Network Installation Management** -> **Perform NIM Administration Tasks** -> **Backup/Restore the NIM Database** -> **Restore the NIM Database**.

Example 6-15 Restore of NIM database using SMIT

```
                  Restore the NIM Database from a Backup

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

                                                      [Entry Fields]
* Filename/Device of the Backup                       [/etc/objrepos/nimdb.backup]



F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F0=Exit             Enter=Do
```

Option 3: Using the command line

Alternatively, you can initiate the restore of the NIM database on your AIX 5L Version 5.2 partition using the command line. In “Saving the NIM database files” on page 119, we used the **tar** command to back up the database. Hence, we need to use this command with different options to restore the database, as shown below:

```
# tar -xvf <device_name>
```

If the host name and network settings on your original NIM master machine and on the new NIM master on the POWER4 partition are the same, you do not need to reconfigure anything. However, it is most likely that these attributes will change. This means that you need to reconfigure the NIM master host name, network adapter address, and make sure that the client names and IP addresses are completed and are correct in /etc/hosts. Additionally, to enable the new NIM master machine to run remotely commands, change the .rhosts file appropriately as well.

We show you in Example 6-16 on page 133 how to change the adapter hardware address using SMIT. The fast path is **smitty nim_chmac**, or just type **smitty** and then select **Software Installation and Maintenance** -> **Network Installation Management** -> **Perform NIM Administration Tasks** -> **Manage Machines** -> **Change / Show Characteristics of a Machine** -> **master** to see the same screen.

Example 6-16 Changing characteristics of NIM machines

```
              Change/Show Characteristics of a Machine

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

[TOP]                                          [Entry Fields]
Machine Name                                   [master]
Hardware Platform Type                         [chrp]                      +
Kernel to use for Network Boot                 [mp]                        +
Machine Type                                    master
Network Install Machine State                   currently running
Network Install Control State                   ready for a NIM opera>
Primary Network Install Interface
Network Name                                    ethernet_1
Host Name                                       [lpar06.itsc.austin.ib>
Network Adapter Hardware Address                [000265af27e8]
Network Adapter Logical Device Name             [ent]
Cable Type                                     N/A                         +
Network Speed Setting                          []                          +
Network Duplex Setting                         []                          +
[MORE...5]



F1=Help              F2=Refresh       F3=Cancel        F4=List
F5=Reset             F6=Command       F7=Edit          F8=Image
F9=Shell             F0=Exit          Enter=Do
```

To verify that all configuration files are restored and will function properly on the new NIM master, do a check as described in Chapter 7, “Post migration tasks” on page 165. Using this method, you restore all NIM database information on your new NIM master, but any other configurations are not transferred. Depending on the complexity of your NIM master, this may or may not be a time-consuming task. Alternatively, you could make a complete backup of your system and restore it on a partition of your POWER4 machine, as described in our third test case.

Case 2: AIX migration and NIM database restore

The second option you have when moving your POWER3 NIM master and its configurations to a POWER4 partition is migrating your original NIM master to AIX 5L and setting up the NIM resources first. Then back up the NIM database as described in the first test case and restore it onto a new installed AIX 5L partition of your POWER4 machine.

For this scenario, our test setup is similar to Figure 6-6. We use the same hardware as before in our first test case, that is, the POWER3 machine running AIX Version 4.3 is an RS/6000 Model F80. The POWER4 partition where we want to restore the migrated NIM database on is an IBM @server pSeries 690 running a new installation of AIX 5L Version 5.2.

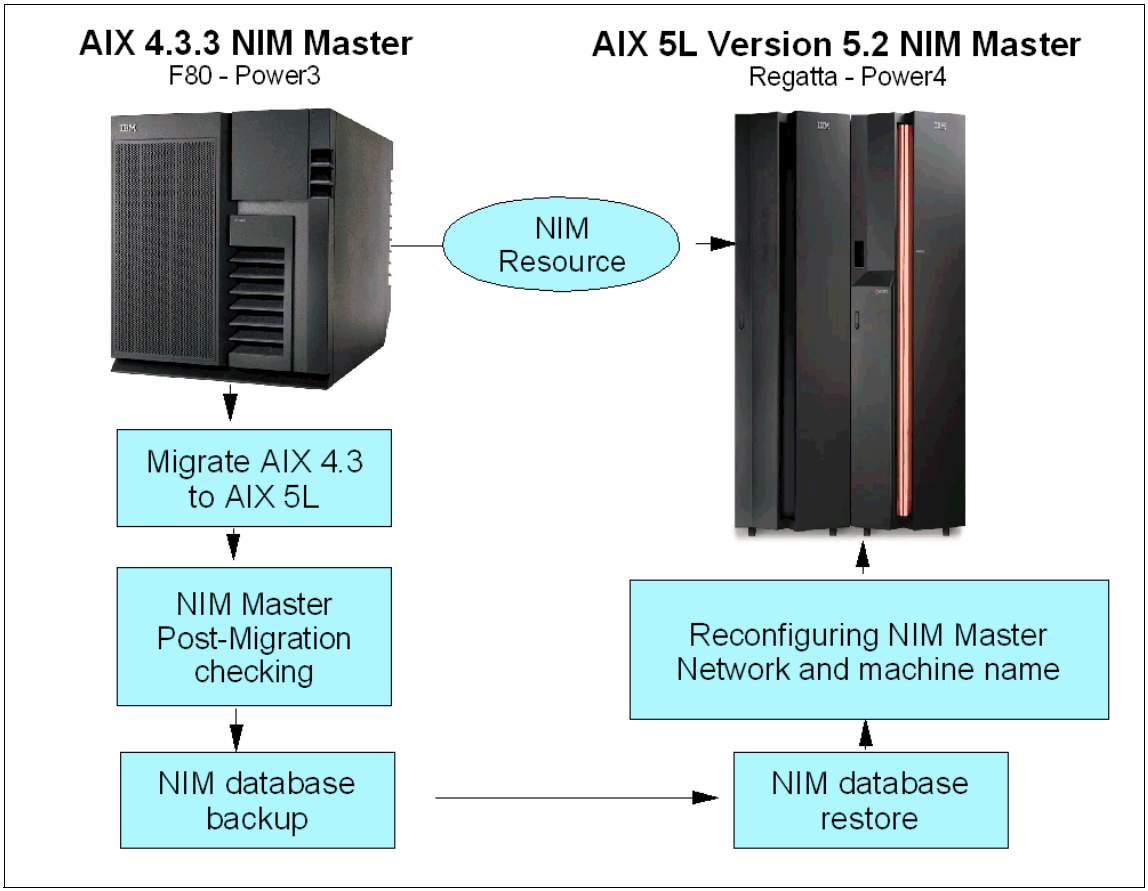

Figure 6-6 Case 2: AIX migration and NIM database restore

The first step in this second test case is that we migrate our NIM master currently running on AIX Version 4.3 to AIX 5L Version 5.2. We perform this operation using media, as described in Chapter 5, “Migration by media” on page 79. We then set up the NIM master with AIX 5L Version 5.2 NIM resources. This way, it can serve as a NIM master for AIX Version 4.3 and AIX 5L Version 5.2. Details of how to configure the NIM master with AIX 5L resources is explained in detail in 6.3, “NIM master migration” on page 117. Having completed all these steps, we can back up the NIM database and restore it onto a new installed partition of our POWER4 machine, that is, an IBM @server pSeries 690. We will not explain this

procedure again, because you have to do exactly the same steps as we described in our first test case.

Case 3: AIX migration and system data restore

If your NIM master has many specific configurations, you might find that test cases 1 and 2 do not appeal to you. The reason for this is that even though you save time by restoring the NIM database rather than setting it up from scratch, you may have many other configurations that do not get transferred. If that is the case, you should make a system image backup, including all NIM resources, and restore it on the POWER4 machine. This way, all your system data and your NIM data will be transferred to the new NIM master, which saves you time.

However, because your NIM master is currently at AIX Version 4.3, you need to first migrate this machine and either set up the NIM master for AIX 5L before the backup or on your new NIM master after the restore. Our test scenario for this third case is shown in Figure 6-7 on page 136. We use exactly the same machine setup as in our two test cases above.

This scenario is very similar to the second case described in 5.3.5, “Case 2: POWER3 to POWER4 AIX 5L Version 5.2 by CD/DVD” on page 95. The only difference is that our system to be migrated and to be moved to a different hardware is a special case because it is the NIM master of our environment. We will not describe the procedure in detail again, but concentrate of the steps involved that are specific to the NIM master.

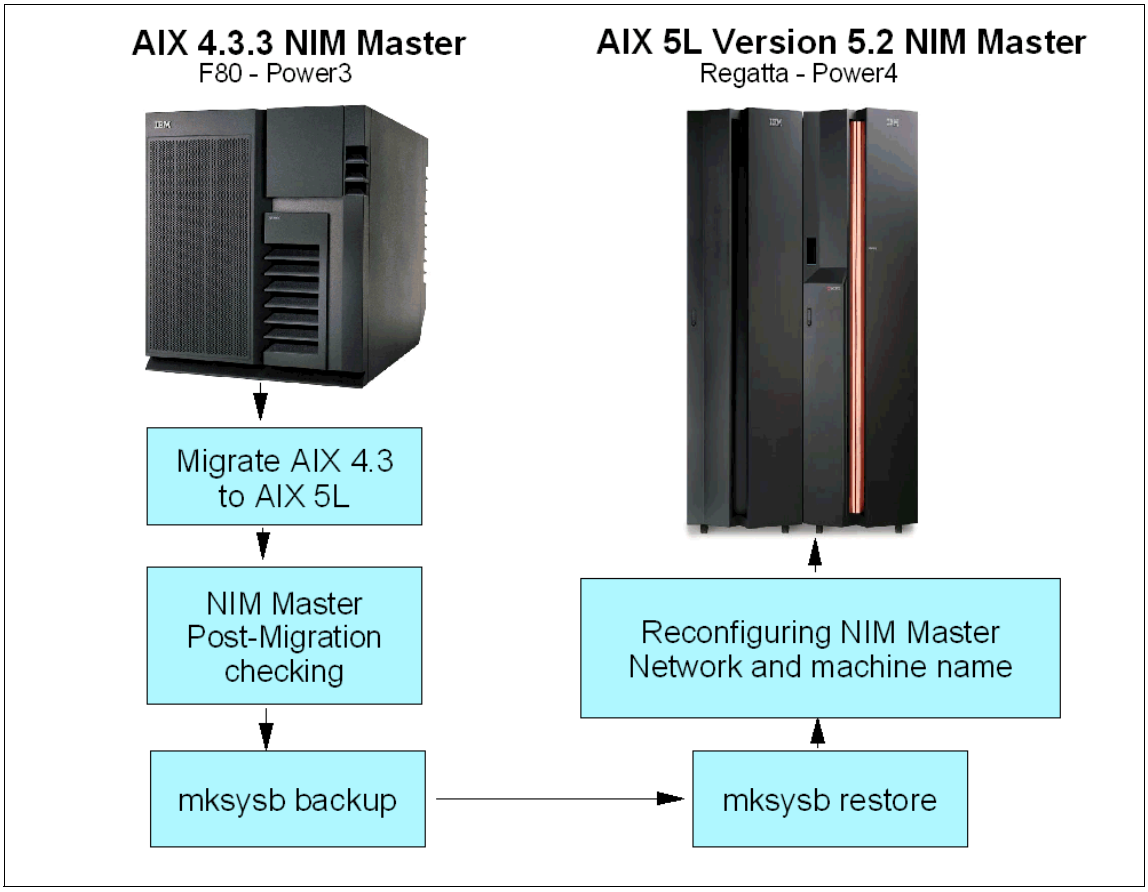

Figure 6-7 Case 3: AIX migration and system data restore

There are two different approaches, which involve the same steps, but in a different order. They are:

Approach 1:

- Migration of AIX Version 4.3 NIM master to AIX 5L
- Configure AIX 5L NIM resources
- Back up migrated and configured NIM master
- Restore on POWER4 partition

Approach 2:

- Migration of AIX Version 4.3 NIM master to AIX 5L
- Back up migrated NIM master
- Restore on POWER4 partition

- Configure AIX 5L NIM resources

Both approaches work and will guarantee you a successful move of your POWER3 NIM master running AIX Version 4.3 to your POWER4 AIX 5L partition.

6.4 NIM client migration

The NIM client migration process, in contrast to the media driven migration or the migration of the NIM master, does not need to manually mount or demount any installation media on the client. Therefore, a high grade of automation during the migration process is provided. This means that multiple migrations can run concurrently. The effort of installing and setting up the NIM master is therefore worth the time and effort. Different methods are possible to get your client machines and their applications onto AIX 5L. These are as follows:

- New installation using NIM
- Alternate disk migration
- Migration using NIM

We discuss the first method in 6.4.2, “NIM installation of client” on page 138 and the NIM migration in 6.4.4, “Migration of NIM client” on page 154. However, a very attractive solution, which we focus on in 6.4.3, “Alternate disk migration of client” on page 140, is the alternate disk migration method. If you have tight requirements on the downtime of your servers, the alternate disk migration will be of great interest to you. The way it works is that you clone your root volume group onto another disk and simultaneously migrate it from AIX Version 4.3 to AIX 5L. The advantage is that during this process, your system is up and running, which means the downtime is reduced to the time it takes to reboot your machine.

6.4.1 Availability during migration

Depending on the method you choose to migrate the operating system, the downtime during the migration process varies depending on whether you use a NIM migration, an alternate disk migration, or a new installation. From the time point of view, a migration takes longer as a new and complete overwrite installation, because it takes more time to upgrade the existing filesets than to just copy them onto a newly installed machine. For both methods, however, your system is down and unavailable for the duration of the migration or installation process. In contrast, for an alternate disk migration, the downtime is reduced to the time it takes to reboot the AIX server.

Furthermore, the fallback time is impacted. This means that there is a significant difference in the time it takes to restore your system in case of a migration or installation failure using your system backup. The advantage of the alternate disk migration method is that you only need to boot from the old root volume group to restore your system, which is left unchanged with this method. This means that the downtime is reduced to the time it takes your server to reboot. We will discuss further advantages and disadvantages of each of these methods in the following sections.

6.4.2 NIM installation of client

One option to bring your NIM client machines to AIX 5L Version 5.2 is to do a new and complete overwrite over the network. The advantage is that you have a freshly installed operating system with default configurations only. Furthermore, the downtime is less than for a migration installation. However, there are disadvantages as well. One time-consuming task is to configure your users and application settings after the installation. To help with this task, you could collect all information prior to the installation and use a script after the installation to restore this information on all client machines.

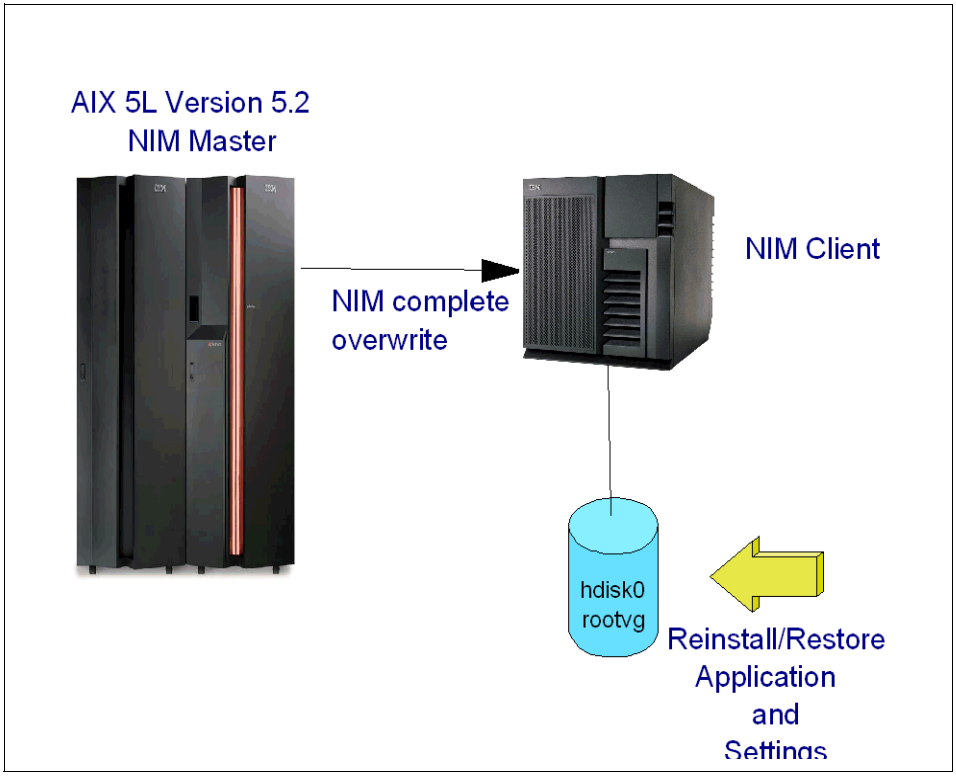

Figure 6-8 New installation of AIX 5L Version 5.2 over the network

We will not discuss the process of a new and complete overwrite installation in more detail. It is well worth mentioning as an option, but we only focus on the migration and alternate disk migration process in the following sections.

However, if you do decide to use the new and complete overwrite method refer to either the *AIX 5L Version 5.1 Installation Guide and Reference*, SC23-4374 or *AIX 5L Version 5.2 Installation Guide and Reference*, SC23-4389 for additional information. One important issue is to define the bosinst.data file for each client on the NIM master. There is only one entry that needs to be adjusted for a new installation of your NIM client. This is shown in Example 6-17 on page 140, where we change the install_method parameter to `overwrite`.

Example 6-17 bosinst.data file for overwrite install

```
control_flow:
    CONSOLE = Default
    INSTALL_METHOD = overwrite
    PROMPT = no
    EXISTING_SYSTEM_OVERWRITE = yes
    INSTALL_X_IF_ADAPTER = yes
    RUN_STARTUP = yes
    RM_INST_ROOTS = no
    ERROR_EXIT =
    CUSTOMIZATION_FILE =
    TCB = no
    INSTALL_TYPE =
    BUNDLES =
    SWITCH_TO_PRODUCT_TAPE =
    RECOVER_DEVICES = Default
    BOSINST_DEBUG = no
    ACCEPT_LICENSES =
    INSTALL_64BIT_KERNEL =
    INSTALL_CONFIGURATION =
    DESKTOP = CDE
    INSTALL_DEVICES_AND_UPDATES = yes
    IMPORT_USER_VGS =
    ENABLE_64BIT_KERNEL = no
    CREATE_JFS2_FS = no
    ALL_DEVICES_KERNELS = yes
    GRAPHICS_BUNDLE = yes
    DOC_SERVICES_BUNDLE = yes
    NETSCAPE_BUNDLE = no
    HTTP_SERVER_BUNDLE = no
    KERBEROS_5_BUNDLE = no
    SERVER_BUNDLE = no
    ALT_DISK_INSTALL_BUNDLE = no
    REMOVE_JAVA_118 = no
```

6.4.3 Alternate disk migration of client

This alternate disk migration method provides the ability to update the filesets of the operating system as if they were normal files in directories. When you start this process, the operating system is copied to another disk on your NIM client. This copied rootvg is then migrated over the network from AIX Version 4.3 to AIX 5L. All this happens while your NIM client is up and running on the original disk. You might have a small performance decrease during the process on your rootvg. As soon as the alternate disk migration is finished, you can reboot to the cloned and migrated disk and run your machine as normal.

This method leads to the shortest downtime for your applications and the operating system. It is a method we recommend for your environment. The main advantages are high flexibility and customization, quick recovery in case of failure, and a reduced downtime. However, with all its advantages, it is also the most complex migration method. This means that what the operating system does automatically is a complex procedure. You as administrator only need to issue one command to start the process.

From a technical point of view, it should be no issue mounting the file systems of the operating system that you want to migrate to another rootvg. The situation you do not want to have is getting duplicate names while there are two rootvgs available in the same environment. There are three issues involved that help to avoid the problem of duplicate names. These are as follows:

- Giving a unique name to the second rootvg
- Using unique alternate mount points for the rootvg file systems
- Ensuring that the migration procedure uses alternate names

This, however, is exactly what is available with the alternate disk migration method, which introduces the **nimadm** command. It means that the naming explained above is nothing you need to worry about. It is part of the alternate disk migration procedure, which can be used with AIX 5L Version 5.2. The new command that enables this method is the **nimadm** command. It will be explained in more detail in our example alternate disk migration in “Alternate disk migration process” on page 149.

The main idea behind the **nimadm** command was to develop a command that you can issue from the NIM master without having to install or to change anything on the NIM client machine. This means that the necessary actions are executed from the NIM master only. For this purpose, the client filesets to be updated are mounted to the NIM master from the cloned rootvg. Even the command to clone the rootvg is executed from the SPOT, which the NIM master exports to the client. The only prerequisite for the client is to be a configured NIM client, which we discuss in “Preparing the NIM client” on page 146.

Furthermore, the migration procedure is designed to work on any AIX Version 4.3 maintenance release levels. This means, there is no need to update your system to the latest maintenance level prior to the migration process. For example, the **alt_disk_install** command executes from the NIM master SPOT. It executes on your AIX Version 4.3 system, even if the SPOT and the **alt_disk_install** command is from AIX 5L Version 5.2.

Overview of method

This first picture of the setup, as displayed in Figure 6-9 on page 142, demonstrates that first part of the alternate disk migration is the cloning of the

root volume group. This execution is initiated from the NIM master to run on the client and to use the SPOT to clone the client's root volume group.

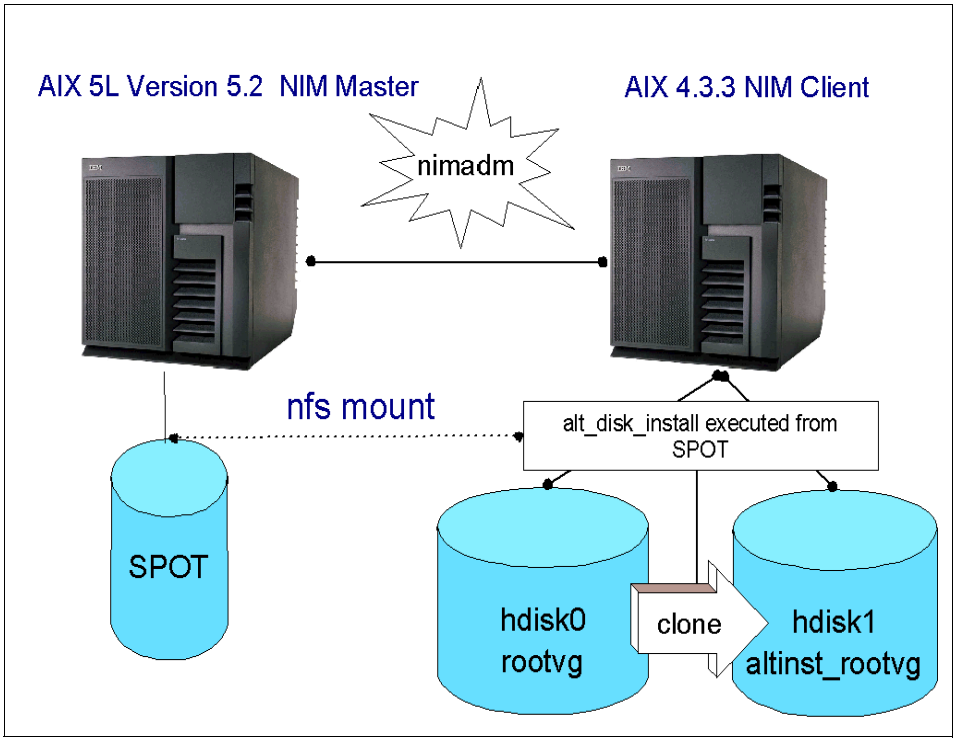

Figure 6-9 Alternate disk migration: The cloning process

The second part of the alternate disk migration process is displayed in Figure 6-10 on page 143. The file systems of the cloned rootvg of the NIM client will get mounted to the NIM master. These will then be updated on the NIM master. For this purpose, the cloned file system names are automatically prefixed with 'alt'. This ensures that they will not conflict with the already mounted original file systems due to duplicate naming.

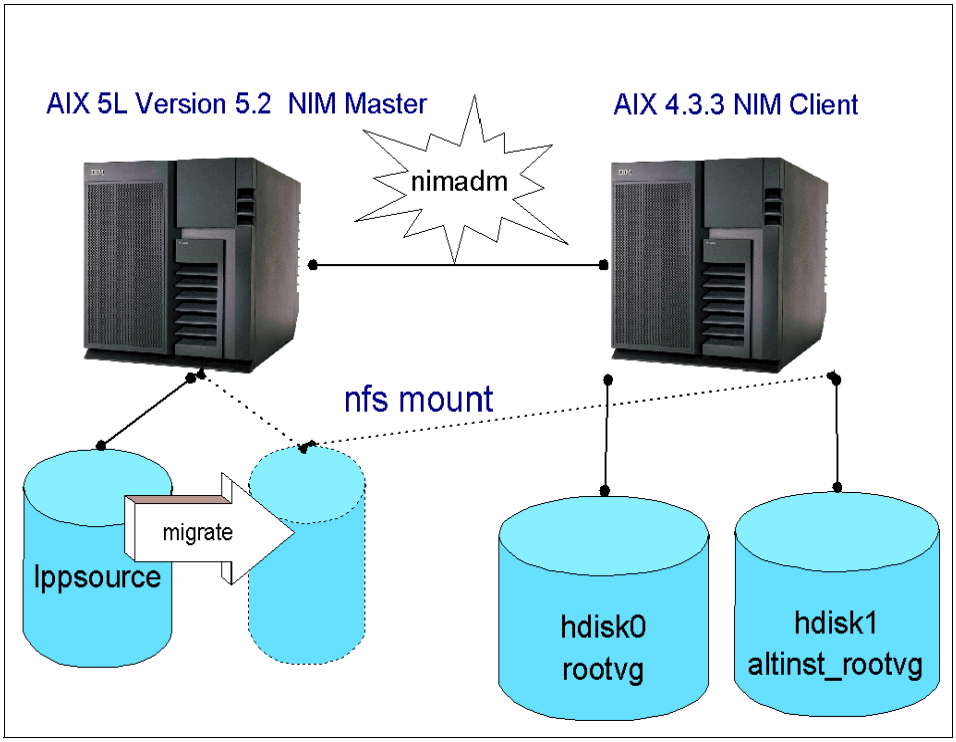

Figure 6-10 Alternate disk migration: The migration process

In Example 6-18, we show a sample output of which file systems of the rootvg are exported during the alternate disk migration process.

Example 6-18 File systems exported during alt_disk_migrate

```
# df -k
Filesystem    1024-blocks      Free %Used    Iused %Iused Mounted on
/dev/hd4            16384      8148   51%     1192    15% /
/dev/hd2           524288     40140   93%    15789    13% /usr
/dev/hd9var         16384     14420   12%      159     4% /var
/dev/hd3            32768     31124    6%       37     1% /tmp
/dev/hd1            16384     15820    4%       20     1% /home
srvr80z.itsc.austin.ibm.com:/nimdata/install/aix52/lppsource/lppsource_aix52
2162688    502640   77%      309     1% /ALT_MIG_IMAGES
srvr80z.itsc.austin.ibm.com:/nimdata/install/aix52/spot_aix52/spot_aix52/usr
360448     14080   97%    10974    13% /ALT_MIG_SPOT
/dev/alt_hd4        49152     37104   25%     1150     5% /alt_inst
/dev/alt_hd1        16384     15820    4%       20     1% /alt_inst/home
/dev/alt_hd10opt       16384     15304    7%       18     1% /alt_inst/opt
/dev/alt_hd3        49152     40668   18%      122     1% /alt_inst/tmp
/dev/alt_hd2       524288     67648   88%    14857    12% /alt_inst/usr
```

```
/dev/alt_hd9var       32768      30300     8%       160      2% /alt_inst/var
```

Example setup

In the following sections, we will show how the alternate disk migration works using our test environment. We have an RS/6000 Model F80 running AIX 5L Version 5.2 as our NIM master machine. It is set up as a NIM master for AIX 5L Version 5.2 using the procedure described in 6.3, “NIM master migration” on page 117.

Our client machine is also an RS/6000 Model F80 running AIX Version 4.3.3 ML 08. This machine consist of six hard disks. Our root volume group is on hdisk0 and hdisk1 and is mirrored as shown in Example 6-19. This is the setup we have before starting the alternate disk migration procedure.

Example 6-19 Volume groups of NIM client before cloning

```
hdisk0          0001813fe67712b5    rootvg
hdisk1          0001813f1a43a54d    rootvg
hdisk2          0001813f95b1b360    None
hdisk3          0001813fc5966b71    None
hdisk4          0001813fc5c48c43     None
hdisk5          0001813fc5c48d8c     None
```

As part of the **nimadm** command, we will clone the rootvg to hdisk2 and hdisk3. These will automatically be named altinst_rootvg. This is shown below in Example 6-20.

Example 6-20 Volume groups of NIM client after cloning

```
hdisk0          0001813fe67712b5    rootvg
hdisk1          0001813f1a43a54d    rootvg
hdisk2          0001813f95b1b360    altinst_rootvg
hdisk3          0001813fc5966b71    altinst_rootvg
hdisk4          0001813fc5c48c43    None
hdisk5          0001813fc5c48d8c    None
```

Prerequisites

As mentioned before, this migration method guarantees the shortest downtime when you migrate your environment from AIX Version 4.3 to AIX 5L, but it also the most complex method and requires a good preparation of your system. You will not be able to go through the whole process using the Web-based System Manager. It only allows you to clone the root volume group. The **nimadm** command for the migration process, however, is not implemented yet. This means that you can use SMIT or the command line to start the alternate disk

migration process. The following prerequisites are necessary for the alternate disk migration method.

Note: You may consider NFS tuning if your system is too slow. Refer to the *AIX 5L Version 5.2 Performance Management Guide*, found at:

http://publib16.boulder.ibm.com/doc_link/en_US/a_doc_lib/aixbman/prftungd/prftungdtfrm.htm

Preparing the NIM master

On the NIM master, you have to install additional filesets to support the alternate disk migration process. These are:

- The latest alternate disk installation filesets
- The latest NIM filesets

A detailed list of filesets including their version number is shown in Example 6-21.

Example 6-21 Prerequisite filesets for NIM master

```
bos.alt_disk_install.rte 5.2.0.0
bos.alt_disk_install.rte 5.2.0.1
bos.sysmgt.nim.master.5.2.0.0
bos.sysmgt.nim.master.5.2.0.1
bos.sysmgt.nim.client.5.2.0.0
bos.sysmgt.nim.client.5.2.0.1
```

If you did not install these filesets, here will not be any menus available in SMIT to clone your root volume group. Trying to start the cloning process using the command line will fail as well, because the client needs to access the alt_disk_install filesets of the SPOT during alternate disk migration.

Important:

- You need to ensure that the prerequisite filesets are installed on your NIM master as well as added to the AIX 5L Version 5.2 SPOT.
- We recommend installing the highest available levels on your NIM master and SPOT. The level of the SPOT and the level of your NIM master must match. You can use 5.2 alt_disk_install utilities to perform a 5.2 or 5.1 migration.

To add missing filesets to your SPOT on the NIM master, use **`smitty nim_task_inst`** as the fast path or just type **`smitty`** and then select **Software Installation and Maintenance** -> **Network Installation Management** ->

Perform NIM Software Installation and Maintenance Tasks -> **Install and Update Software**.

Select **SPOT** as resource type you want to add software to and the appropriate filesets.

Preparing the NIM client

You do not need any additional filesets on the client. Your NIM client needs to be configured as a client machine, for which you need the bos.sysmgt.nim.client fileset. If your machine was a NIM client before, the fileset will all be set up already. The **nimadm** command then executes on the NIM master. The **alt_disk_install** command runs on the client, but is executed from the SPOT, which is mounted from the NIM master.

Furthermore, the NIM client must be accessed from the NIM master via a remote shell (rsh). To ensure that this works, there must be an .rhost file on the client machine that contains the following lines. The first line demonstrates the general syntax. The second line is what it looks like on our test machine:

```
<hostname of nim master>    root
srvr80j                    root
```

For security reasons, root must be the user name behind the host name. Additionally, only the root user may be permitted to read-write the .rhosts file. Otherwise, the rsh will not work. The correct permissions are shown below:

```
-rw------- 1 root system 13 Jan. 24 15>08 .rhosts
```

If the NIM client is not already set up to be a client, you must do it in the next step. This can be done by using the Web-based System Manager, SMIT, or command line. We will only demonstrate this setup procedure using SMIT. Type the **smitty niminit** fast path or just type **smitty** and then select **Software Installation and Maintenance** -> **Network Installation Management** -> **Configure Network Installation Management Client Fileset** to get the screen shown in Example 6-22 on page 147. Fill in the client's host name, the network install interface, the NIM master host name, and a few others, as shown in Example 6-22 on page 147.

Example 6-22 Configuring a NIM client

```
        Configure Network Installation Management Client Fileset

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

                                                      [Entry Fields]
Machine Name                                          [srvr80j]
Primary Network Install Interface                     [en1]        +
Host Name of Network Install Master                   [srvr80z]

Hardware Platform Type                                chrp
Kernel to use for Network Boot                        [mp]         +
Comments                                              []

Alternate Port Numbers for Network Communications
(reserved values will be used if left blank)
Client Registration                                   []           #
Client Communications                                 []           #



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

Having completed the configuration of the NIM client task, you will now have an /etc/niminfo file on your client machine. This is necessary, because the NIM master needs to access this file on the client machine during the alternate disk migration process. Example 6-23 on page 148 shows the /etc/niminfo file of our test client machine with host name srvr80j. As you can see, the NIM master machine in our environment is called srvr80z.itsc.austin.ibm.com.

Example 6-23 Example of /etc/niminfo file of NIM client

```
                    Network Install Manager
warning - this file contains NIM configuration information
          and should only be updated by NIM

export NIM_NAME=srvr80j
export NIM_HOSTNAME=srvr80j.itsc.austin.ibm.com
export NIM_CONFIGURATION=standalone
export NIM_MASTER_HOSTNAME=srvr80z.itsc.austin.ibm.com
export NIM_MASTER_PORT=1058
export NIM_REGISTRATION_PORT=1059
export NIM_BOS_IMAGE=/SPOT/usr/sys/inst.images/installp/ppc/bos
export NIM_BOS_FORMAT=rte
export NIM_HOSTS=" 9.3.4.35:srvr80j.itsc.austin.ibm.com
9.3.4.33:srvr80z.itsc.austin.ibm.com "
export NIM_MOUNTS=""
```

This file enables you to also monitor the state of the client machine during the alternate disk migration process. You can do this by typing the **`lsnim -l <client name>`** command on the NIM master machine.

Another concern is the Trusted Computing Base (TCB). As already mentioned in Chapter 4, “Planning and preparation” on page 37, if you want to have the TCB enabled after the migration, you cannot use alternate disk migration. You need to either migrate your system by media, use the alternate disk migration, or use the network migration with TCB disabled with AIX 5L Version 5.2.

Attention: You must disable TCB if you want to use alternate disk migration.

If you decide to use the second option, which means TCB enablement is no longer needed, it is possible to disable the Trusted Computing Base using the following commands. First, you need to get the Object Data Manager (ODM) content to a file. Type the following:

```
odmget -q attribute=TCB_STATE PdAt > own_filename
```

In this own_filename, you need to edit one line, as shown in Example 6-24 on page 149. Change `tcb_enabled` to `tcb_disabled` and save the file.

Example 6-24 Disable Trusted Computing Base

```
PdAt:
        uniquetype = ""
        attribute = "TCB_STATE"
        deflt = "tcb_enabled"
        values = ""
        width = ""
        type = ""
        generic = ""
        rep = ""
        nls_index = 0
```

To change these values in the ODM, type the following two commands:

```
odmdelete -o PdAt -q attribute=TCB_STATE
odmadd own_filename
```

The Trusted Computing Base is no longer enabled. To verify this situation, type the **odmget** command.

Alternate disk migration process

To start the alternate disk migration of your NIM client's rootvg, you have two options: SMIT and the command line on your NIM master machine. At the time of writing, there are no menus for alternate disk migration available with Web-based System Manager.

Note: NFS tuning is required to optimize performance during the alternate disk migration process. Refer to the AIX product documentation for detailed information on tuning parameters.

Option 1: Using SMIT

When you are using SMIT, you can get to the menu shown in Example 6-25 on page 150 either by issuing the **smitty nimadm** or by typing **smitty** and then selecting **Software Installation and Maintenance** -> **Alternate Disk Installation** -> **NIM Alternate Disk Migration** -> **Perform NIM Alternate Disk Migration**. However, remember that this command is issued at the NIM master machine, *not* on the NIM client.

Example 6-25 Initiate alternate disk migration using SMIT

```
                Perform NIM Alternate Disk Migration

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

                                                      [Entry Fields]
* Target NIM Client                                   [srvr80j]               +
* NIM LPP_SOURCE resource                             [lppsource_aix52]       +
* NIM SPOT resource                                   [spot_aix52]            +
* Target Disk(s) to install                          [hdisk2 hdisk3]

  NIM IMAGE_DATA resource                             []                      +
  NIM BOSINST_DATA resource                           []                      +
  NIM EXCLUDE_FILES resource                          []                      +
  NIM INSTALLP_BUNDLE resource                        []                      +
  NIM PRE-MIGRATION SCRIPT resource                   []                      +
  NIM POST-MIGRATION SCRIPT resource                  []                      +

  Phase to execute                                    [all]                   +
  NFS mounting options                                []
  Set Client bootlist to alternate disk?               yes                    +
  Reboot NIM Client when complete?                     no                     +
  Verbose output?                                      no                     +
  Debug output?                                        no                     +

  ACCEPT new license agreements?                       yes                    +


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

There are optional and required entries for this menu. We will focus only on the required input. This includes the name of the NIM client machine. In our test case, this is srvr80j. The **nimadm** command also supports migrating several clients simultaneously. You need to specify the SPOT and lpp_source name, which are stored on the NIM master machine. We call these resources spot_aix52 and lppsource_aix52, respectively. You need to specify the disk or disks you want to clone and migrate your rootvg to. In our test case, we have the rootvg on hdisk0 and hdisk1 and choose hdisk2 and hdisk3 for the clone. Furthermore, it is important to set the Accept New License Agreement to Yes.

In the above example, we also specified to have all phases executed. You could also specify to have only certain phases done. We will explain in more detail what operation the system performs during which phase. In general, you will

want to execute all phases, which is the default value. One option is to specify whether or not the system is automatically rebooted after the process is finished, and put the cloneds or migration system into production. We chose not to have an automatic reboot after completion of the alternate disk migration. The reason for this is that we have applications running on our rootvg when we perform the alternate disk migration. This means that only when we are sure this process is completed successfully, we will stop the applications and initiate the reboot and a change of the bootlist manually. Which option suits you best depends on your environment.

In Chapter 5, “Migration by media” on page 79, we describe the AIX migration by media. We explain all possible install options and mention that at this stage, we are not able to select additional software to be installed. This is not a problem, because we can do it manually after the migration process. However, with the alternate disk migration menu shown in Example 6-25 on page 150, we can specify resources to be excluded or to be installed simultaneously. These include BOSINST_DATA, IMAGE_DATA, EXCLUDE_FILES, and INSTALLP_BUNDLE resources.

Option 2: Using the command line

The second option to start the alternate disk migration process of your NIM client from AIX Version 4.3 to AIX 5L Version 5.2 by issuing the following command on the command line:

```
nimadm -c srvr80j -s spot_aix52 -l lppsource_aix52 -d “hdisk2 hdisk3” -Y
```

As before, you specify the NIM client name, the lpp_source and SPOT resource, the target disks, and, very important, you accept the license agreement by using the -Y flag. This is required for some filesets. Your system might hang when you forget to set the -Y flag. There are additional flags available to specify all the details, as mentioned with the SMIT menu. For further information, refer to the AIX 5L Version 5.2 documentation.

When you initiate the alternate disk migration process on your NIM master, you will be informed about which phase of the **nimadm** command is executed at the moment and some more information. There is no further step for you to do at this time unless the process fails. However, we do want you to have some knowledge about the procedure behind this command, which we explain in the following section.

The nimadm command

When you initiate the **nimadm** command on your NIM master machine, you start a process that involves 12 phases. It is done by the operating system automatically and involves two major parts: cloning the rootvg and migrating this clone to AIX 5L Version 5.2.

The performed steps or phases of the **nimadm** command are as follows:

- Phase 1: The **alt_disk_install** command is issued from the NIM master to the client. This clones the rootvg of the NIM client to another spare disk or disks on the client machine.
- Phase 2: The NIM master runs remotely the **export** command on the client machine. This way, the file systems on the cloned rootvg disk of the client are exported.
- Phase 3: The exported client file systems are NFS mounted by the NIM master.
- Phase 4: Customized pre-migration scripts will be executed if specified.
- Phase 5: System configuration files are saved. Migration space is calculated. File systems get expanded.
- Phase 6: All system filesets of the cloned NIM client disk will be migrated using the **installp** command, including required RPM images.
- Phase 7: Customized post-migration scripts will be executed if specified.
- Phase 8: The boot image of the migrated alternate rootvg is written to the boot record using the **bosboot** command.
- Phase 9: All mounts made on the NIM master are removed.
- Phase 10: All exports of the clients file systems are removed.
- Phase 11: The **alt_disk_install** command is used again to make final adjustments. It puts the altinst_rootvg to sleep. The bootlist is set to the target disk.
- Phase 12: Cleanup tasks are performed. The reboot of the NIM client is initiated, if specified.

For a detailed description of each phase and the **nimadm** command in general, refer to the *AIX 5L Version 5.2 Commands Reference*, found at:

`http://publib16.boulder.ibm.com/cgi-bin/ds_rslt#1`

If you want to view the log files for each client, go to /var/adm/ras/alt_mig/<NIM_client_name>_alt_mig.log on the NIM master machine.

In Example B-3 on page 224, we show you a sample output that you will either find in the mentioned log files or which you obtain when starting the alternate disk migration process using SMIT or the command line. The phases are listed as your process runs. This ensures that you understand, at each time of the migration process, what operations the machine performs. Additionally, if you run into problems, you can easily verify at which phase of the process you encountered the problems. This makes troubleshooting a lot easier.

Restarting the migration process

If for some reason the migration process breaks up, the cleanup is done automatically by the operating system. All you need to do is try and analyze what has gone wrong and then reissue the **nimadm** command either using SMIT or the command line. However, during our test migrations, we had some situations where the cleanup was not done properly. For this task, you can explicitly call the command below to do only the cleanup procedure again:

```
nimadm -C srvr80j -s spot_aix52 [-F]
```

The following examples explain some tasks you may have to do to clean up the unsuccessful alternate disk migration process. These include the following:

- Unmount all /dev/alt_* file systems:

  ```
  /dev/alt_hd4
  /dev/alt_hd1
  /dev/alt_hd10opt
  /dev/alt_hd3
  /dev/alt_2
  /dev/alt_hd9var
  ```

- Varyoff the volume group altinst_rootvg\.
- Export the altinst_rootvg using the **alt_disk_install -x** command.
- Remove the /dev/alt_* file systems.

The recommendation described in the README file is to avoid using the **exportvg** command to clean up the altinst_rootvg. Instead, you should use the **alt_disk_install -x** command. The reason is that with the **exportvg** command some file system stanzas in /etc/filesystems get deleted. Even though we did not encounter this problem, it is best to follow the recommendation to ensure your system will function as desired.

Note: To clean up your altinst_rootvg, you should use the **alt_disk_install -x** command rather than the **exportvg** command for the reasons described in the preceding paragraph.

Furthermore, you need to verify if you can start and stop NFS daemon. Activate the NFS daemon to remove the remaining file systems. If you fail to remove all the exports, the alternate disk migration process will not run successfully when you next initiate it. You will see the following error message at Phase 2:

```
mknfsexp: 1831-355 an export for <filesystemname> already exists
0505-154 nimadm: Error exporting client alt_inst filesystems.
```

When you issue the **nimadm** command again, all twelve phases will start and run through. This includes cloning the root volume group and migrating the clone to

AIX 5L Version 5.2. You do not have to manually clone your system because it is done automatically by the **nimadm** command. You can refer to 5.3.6, “Case 3: Cloning your system” on page 100 for more information on the cloning process and how to initiate it manually. There are two ways to initiate cloning when you are in an NIM environment: From the NIM master and from the NIM client. The process itself and the involved SMIT or Web-based System Manager menus are the same as we describe in 5.3.6, “Case 3: Cloning your system” on page 100.

6.4.4 Migration of NIM client

The migration method via NIM is similar to the migration via CD with one difference: that all migration activities take place over the network connection between the NIM master and the machines being migrated, which are called NIM clients. In the case of the NIM migration, the client gets its boot record via the bootp protocol from the NIM master. With this boot record, the NIM client creates a RAM file system. This allows basic UNIX operations that are necessary for installation to take place, such as NFS mounts or trivial file transfer protocol (tftp). The migration itself is running on the client side with a data connection to the NIM master. This means that all files that are migrated and their attributes are fully accessible.

Compared to the other two methods, that is NIM new and complete overwrite installation and alternate disk migration installation, the migration over the network is the one that leads to the longest application downtime if things go wrong and you need to restore the operating system. On the other hand, it is the most straightforward option, which involves the least effort for the administrator if you plan and prepare it well so that it works without errors. Most configuration parameters will be available after migration, including users and passwords.

Migrating the NIM client over the network means that we boot the AIX machine over the network with a network boot record. This uses NFS mounts to update the filesets on the client machine. The migration process runs on the target machine and does not allow any application to be active. The reason for that is that the booted kernel runs in single user mode and only supports the NIM required actions. The filesets being migrated are not backed up during this process, which means you have to keep your own backup. In case of a migration failure or any other errors, you need to restore the backup of the rootvg.

The setup we use in our test environment is shown in Figure 6-11 on page 155. The NIM master is an RS/6000 Model F80 running AIX 5L Version 5.2. This machine is set up as a NIM master with AIX 5L Version 5.2, as described in 6.3, “NIM master migration” on page 117. The NIM client machine is also an RS/6000 Model F80 running AIX Version 4.3.3 ML08. The two machines are connected via Ethernet.

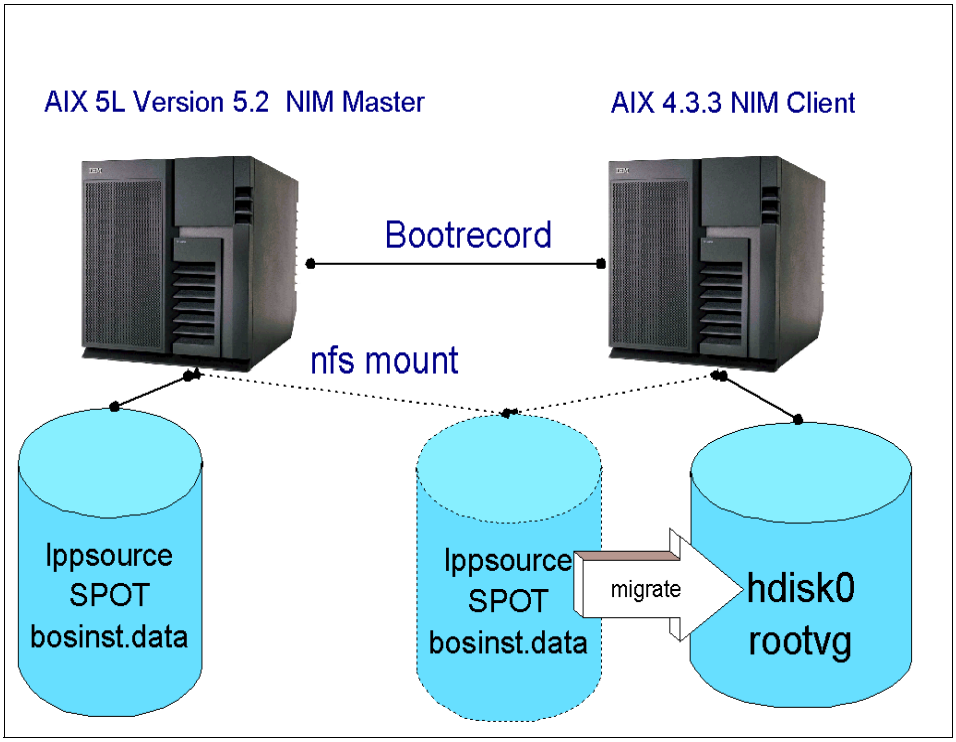

Figure 6-11 NIM migration

Alternatively, if for any reason you cannot use the alternate disk migration method because of the TCB restrictions, for example, you can use the NIM migration in conjunction with cloning the rootvg. This means that you clone the rootvg manually using the **`alt_disk_install`** command, as described in Chapter 5, “Migration by media” on page 79, and then migrate the cloned rootvg. You have to boot over the network again, which means you can limit the downtime. However, the advantage is that you keep your original rootvg available. If things go wrong, all it takes is a reboot to this rootvg and the system is up and running again. That saves you having to restore your mksysb backup as well. Out test setup is shown in Figure 6-12 on page 156. We tested this variation of the migration method successfully.

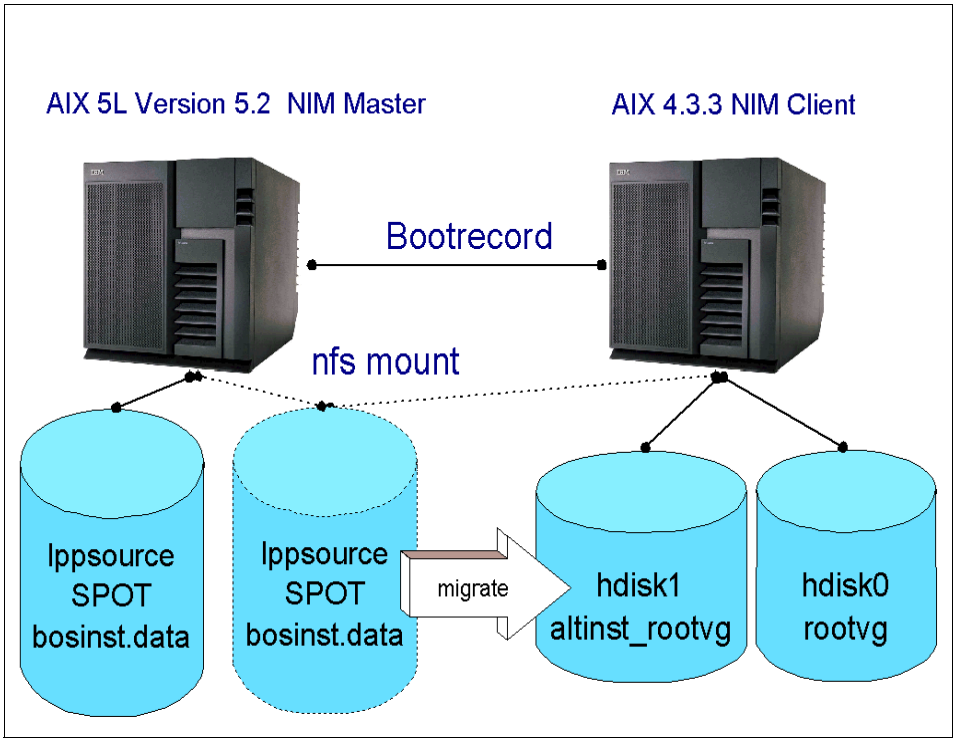

Figure 6-12 Migration of cloned rootvg

Prerequisites

One important requirement is that the network card of the NIM client machine supports the bootp protocol. A workaround, if you have no supported bootp protocol, is to create a floppy disk that extends the network card. This is a feature included in the NIM software package. This means that whichever way you use, your AIX machine should be able to boot and install over the network, so you will be able to use all the advantages of Network Installation Management.

There are some requirements for the NIM master that must be met before you can start the client migration. These are only mentioned in this section, but will be explained in detail in 6.3, “NIM master migration” on page 117. They include that the NIM master must be on a AIX release level equal to or higher than the client machine. Furthermore, there must be some resources available on the NIM master that the client uses during the migration process. If we assume that we want to migrate to AIX 5L Version 5.2, these are:

- AIX 5L Version 5.2 lpp_source
- AIX 5L Version 5.2 SPOT resource
- Bosinst.data file

If there is not already a bosinst.data file available on the NIM master, a template can be obtained from the /usr/lpp/bosinst directory. Copy the bosinst.template file to a location of your choice and modify it according to your needs. Then define it as a NIM resource. In Example 6-26, we show an example of the modified bosinst.data file. The ACCEPT_LICENSES stanza specifies whether to accept software license agreements during the BOS installation. The default value is `no`. To automatically accept them, set this to `yes`.

Example 6-26 Sample bosinst.data file

```
control_flow:
    CONSOLE = Default
    INSTALL_METHOD = migrate
    PROMPT = no
    EXISTING_SYSTEM_OVERWRITE = yes
    INSTALL_X_IF_ADAPTER = yes
    RUN_STARTUP = yes
    RM_INST_ROOTS = no
    ERROR_EXIT =
    CUSTOMIZATION_FILE =
    TCB = no
    INSTALL_TYPE =
    BUNDLES =
    SWITCH_TO_PRODUCT_TAPE =
    RECOVER_DEVICES = Default
    BOSINST_DEBUG = no
    ACCEPT_LICENSES =
    INSTALL_64BIT_KERNEL =
    INSTALL_CONFIGURATION =
    DESKTOP = CDE
    INSTALL_DEVICES_AND_UPDATES = yes
    IMPORT_USER_VGS =
    ENABLE_64BIT_KERNEL = no
    CREATE_JFS2_FS = no
    ALL_DEVICES_KERNELS = yes
    GRAPHICS_BUNDLE = yes
    DOC_SERVICES_BUNDLE = yes
    NETSCAPE_BUNDLE = no
    HTTP_SERVER_BUNDLE = no
    KERBEROS_5_BUNDLE = no
    SERVER_BUNDLE = no
    ALT_DISK_INSTALL_BUNDLE = no
    REMOVE_JAVA_118 = no

target_disk_data:
    PVID = 00015f9f9ec895a4
    SAN_DISKID = none
    CONNECTION = scsi0//2,0
    LOCATION = 11-08-00-2,0
```

```
    SIZE_MB = 8678
    HDISKNAME = hdisk0

target_disk_data:
    PVID = 00015f9f3919bbf1
    SAN_DISKID = none
    CONNECTION = scsi0//4,0
    LOCATION = 11-08-00-4,0
    SIZE_MB = 8678
    HDISKNAME = hdisk1

locale:
    BOSINST_LANG = en_US
    CULTURAL_CONVENTION = en_US
    MESSAGES = en_US
    KEYBOARD = en_US
```

For the migration of the NIM client, you need sufficient disk space, just as with the migration by CD. If you do not have enough disk space available, the migration process will run into several problems. This usually means you will have to restore the rootvg backup and start the migration again. In Chapter 4, “Planning and preparation” on page 37, we list the disk space needed when you want to migrate from AIX Version 4.3 to AIX 5L. Additionally, in Example 6-27, we show the usage of the rootvg file systems before and after the migration to AIX 5L Version 5.2.

Example 6-27 File system size of rootvg before and after migration

```
Before migration:
# df -kI
Filesystem    1024-blocks      Used      Free %Used Mounted on
/dev/hd4            16384      7108      9276   44% /
/dev/hd2          1048576    495200    553376   48% /usr
/dev/hd9var         16384      5596     10788   35% /var
/dev/hd3            32768      1744     31024    6% /tmp
/dev/hd1            16384       564     15820    4% /home

After migration:
# df -kI
Filesystem    1024-blocks      Used      Free %Used Mounted on
/dev/hd4            16384     11364      5020   70% /
/dev/hd2          1048576    655112    393464   63% /usr
/dev/hd9var         16384      9052      7332   56% /var
/dev/hd3            32768      1292     31476    4% /tmp
/dev/hd1            16384       564     15820    4% /home
/proc                   -         -         -    -  /proc
/dev/hd10opt        32768      8000     24768   25% /opt
```

Preparing the NIM master for the migration process

The NIM master must be ready to serve the client's boot request when the migration process is initiated. This means, you need to allocate the appropriate resources to your client and make the NIM master aware of the type of operation to perform when receiving a client's boot request. There are three ways to prepare the NIM master for the migration process. These include the Web-based System Manager, SMIT, and command line. However, we will only outline the steps to perform, because this is exactly the same as you would do when you want to install a NIM client. For additional information, refer to the publication *AIX 5L Network Installation Management Guide and References* SC23-4385.

For the main resource allocation, you must select the following:

- Shared Product Object Tree (SPOT)
- Software source (lpp_source)
- Bosinst.data resource

You can start the migration process from the NIM master or from the NIM client. The procedure is exactly the same as for a NIM installation. Instead of choosing **New and Complete Overwrite** you will select **Migration** as the installation type. The selection between installing and migrating is already done automatically if you put `INSTALL_Method = migrate` in the bosinst.data file, as described in Example 6-26 on page 157. Furthermore, you should check the **Accept new license agreement** box at the bottom of the screen.

Booting the NIM client

There are three ways to initiate a netboot of the NIM client to begin the migration process. These include:

- Set the network adapter as the first boot device on the NIM client and reboot the system.
- Boot into SMS menu and configure the network adapter as the first boot device.
- Use the SMIT menu **Perform a Network Install** from the client.

During the reboot, the machine will get its boot record according to the bootlist from the NIM master over the network. The client distributes a broadcast and the NIM server that has stored the NIM client's MAC address of the broadcasting adapter will send back the boot record. The NIM master knows which client to serve because you have allocated the resources. Refer to 6.4.5, "Troubleshooting" on page 161 if your NIM client has problems booting over the network.

Migration process

If the NIM client machine boots over the network as desired and all options in the bosinst.data file are set according to the migration needs, the NIM client will work on the boot record. It will mount the necessary files from the NIM master and it display `migration preparation menu`. This gives you reassurance that you are doing a migration and not an overwrite installation. At a later time, another line will appear that states `bos installation menu`. This is a critical time for your system, because you have lost your old operating system at this stage. If any errors occur, the only way to get your operating system back and running is by restoring the backup you made prior to the migration process.

Further activities of the migration will be displayed on the screen. This starts with the file system allocation and a listing of which filesets are being updated. Additionally, a log is created in the client's /var/adm/ras directory, which is called devinst.log. Furthermore, the current status is reported to the NIM master and can be queried using the **`lsnim -l <clientname>`** command on the NIM master machine.

The migration is finished when you get a message saying `overmounting ./`, which means the RAM file system is overmounted by the now migrated original /(root) file system. Final migration activities like writing the boot record and the bootlist are then taking place. The NIM server will be informed of the successfully completed migration and resets the NIM clients state. The NIM master ensures that the client does not boot over the network again.

The final step will be the reboot of your NIM client machine. It will then present itself as being migrated to the new operating system release of AIX 5L Version 5.2. We recommend you make another backup right after the migration process before you start to do any post migration tasks. In addition, in 6.4.5, "Troubleshooting" on page 161, we give some ideas of possible problem situations as well as how to fix and avoid these problems.

Restarting the migration process

Generally, the NIM master is automatically informed by the migration process of a successful migration. The NIM server resets the state of the client to prevent an endless loop of network boots. The resources allocated to the client will be deallocated.

If, for any reason, your NIM migration fails, you have to do some additional tasks to reset the NIM environment. This usually means that the system hangs. The NIM master cannot reset the state of the client or deallocate its resources. You need to do these tasks manually.

To reset the state of the NIM client using SMIT, select the **`smitty nim_mac_op`** fast path on the NIM master. The screen shown in Example 6-28 on page 161

will appear. Select your client machine name as well as yes for deallocation of resources and yes for the Force operation. The Force operation is designed to prevent unwanted resets of clients during normal installation operations.

Example 6-28 Resetting the NIM state of a client after an unsuccessful migration

```
                  Reset the NIM State of a Machine

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

                                                      [Entry Fields]
Target Name                                           srvr80j
Deallocate All Resources?                             yes         +
(removes all root data for
diskless/dataless clients)
Force                                                 yes         +



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

After resetting the state of the client, the NIM master is ready to perform new operations on that client again. You need to restore your backup of the rootvg to be able to work with the client machine again. If you have processed the migration on a previously cloned disk, you only have to reboot your machine to your original rootvg. Repair the problems that caused the migration to fail and start the migration process again.

6.4.5 Troubleshooting

Even though all our test migrations worked successfully, you might still run into problems. For this reason, we include some useful hints and tips in this section that will help you identify the source of the problem that causes the migration to fail.

Log files

The first place to look if your migration process run into problems are the log files. The information in there should help you investigate the reason for the problems. You can do this with the Web-based System Manager by right-clicking on the machine name, as shown in Figure 6-13 on page 162.

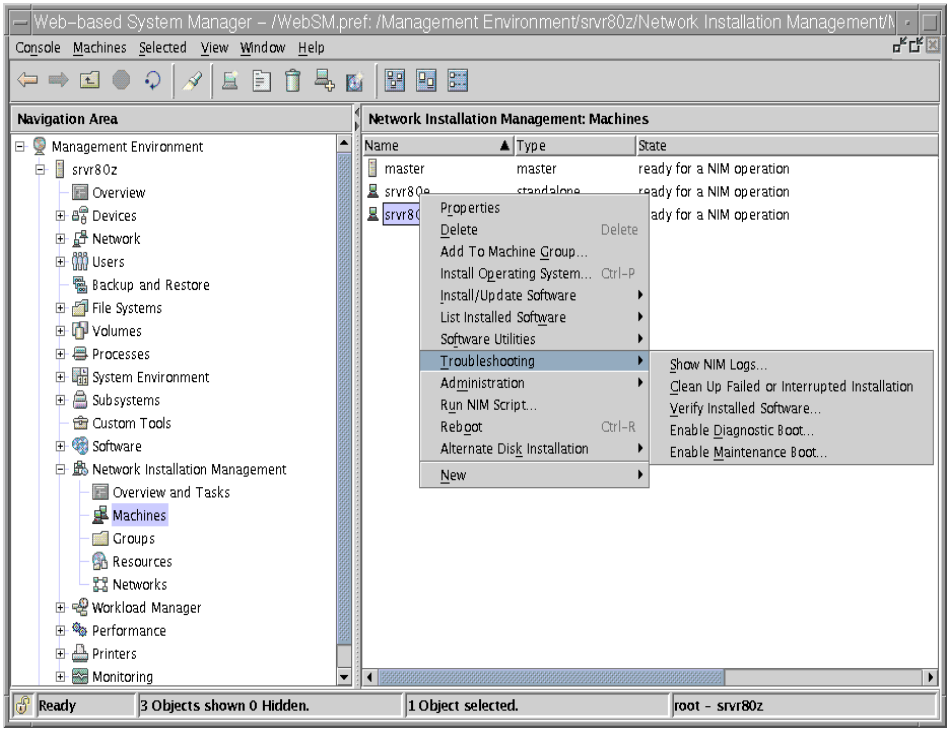

Figure 6-13 Migration log files

There is a variety of log files you can look at that are displayed in Figure 6-14 on page 163. You can obtain the same information by directly viewing the log files.

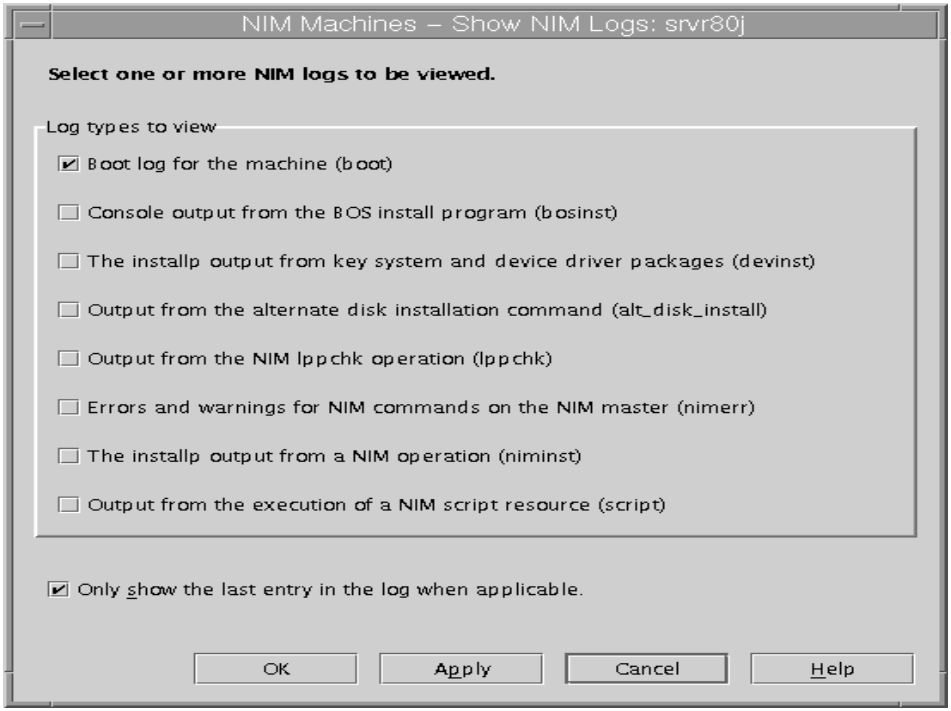

Figure 6-14 Available log files after the migration process

Booting over the network

When using NIM, one of the most common problems has to do with booting over the network. If your machine has problems getting a boot record, check the following issues, which may be the cause of your problems.

- The client's SMS menu already has IP addresses, netmask, and/or gateways defined that do not match the current environment. You can test the connection between NIM master and client using the ping function as part of the SMS menu. For this test, you need to put in entries for IP addresses.However, after the successful testing, it is best to set all fields to zero again. This way, the values stored in the NIM master are used.
- The adapter speed is a critical parameter. If you have the adapter speed set to auto on your NIM client and master, you might want to change it. It usually does not work as required. We recommend you set the adapter speed and duplex value to a desired value on both machines.
- Check the microcode level of your machine and of your adapters. An old microcode level is often the cause for network problems between NIM master and client. You might not even notice it during normal business on your AIX

machine, but the network boot needs the successful initializing of the network card.

Migration hangs during the migration process

If the migration has started and gets the boot record, a possible error situation is that it hangs at a later time and does not seem to do anything. Check one of the following issues:

- Are the file systems properly exported?
- Can the client access the file systems of the NIM master via NFS?

Furthermore, we encountered problems when migrating mirrored rootvgs. Using the same procedure without the mirror worked without any problems. This is a known problem. The only workarounds at the time of the writing of this redbook are to unmirror the disks before the migration and mirror them again afterwards or to use alternate disk migration. This method works fine with mirrored disks.

7

Post migration tasks

Before you can use your migrated system in a production environment again, you need to complete some post migration tasks. These will be described in detail in the following sections. We first concentrate on general tasks that you need to complete for any system no matter whether it is a NIM master or client or whether you migrate your system using media. In addition, we include information on NIM master specific tasks and post migration issues for the alternate disk migration method.

Furthermore, we will concentrate on the most important post migration tasks that will affect the majority of migrated AIX machines only. This means that if you have a complex environment with a large number of servers that need to be migrated from AIX Version 4.3 to AIX 5L, you may consider doing a test migration. Using the information obtained, you could create post migration scripts that will keep the post migration effort and time frame down to a minimum. The checks could then be very specific to your environment in addition to the general outlined topics below which we will discuss in this chapter:

- Installation assistant
- Merging of configuration files
- Network and security aspects
- Mirrored disks

7.1 Common tasks

There are several post migration tasks that must be completed before your system is fully functioning again. These differ depending on which migration method you use. Some tasks are identical, while others are specific to a certain migration method. We concentrate on the post migration issues that are common to all migrated systems.

Installation assistant

After each migration installation and the automatic reboot of your server, the Installation Assistant Menu needs to be completed. It will ask you to do some basic configuration tasks as outlined in detail in the AIX 5L Version 5.1 Installation Guide and Reference, SC23-4374 or AIX 5L Version 5.2 Installation Guide and Reference, SC23-4389. On a graphical console, you will see the installation assistant as a graphical interface. On a non-graphical console, it is an ASCII display.

When doing a migration installation by media, the installation assistant automatically starts after reboot. When you migrate your system using Network Installation Management (NIM), you can select whether or not you want this task menu to run automatically after reboot. However, even if you disable this function to run automatically, you can start it manually at any time after the migration process is finished. You can do so either by using the **`smitty assist`** fast path or the **`install_assist`** command.

The most important task to complete is to accept the license agreement. The next window is shown in Figure 7-1 on page 167. You need to set and verify the time and date on your AIX server and the root password, as well as manage paging space and the TCP/IP configuration. The menu looks just as it does after an operating system installation.

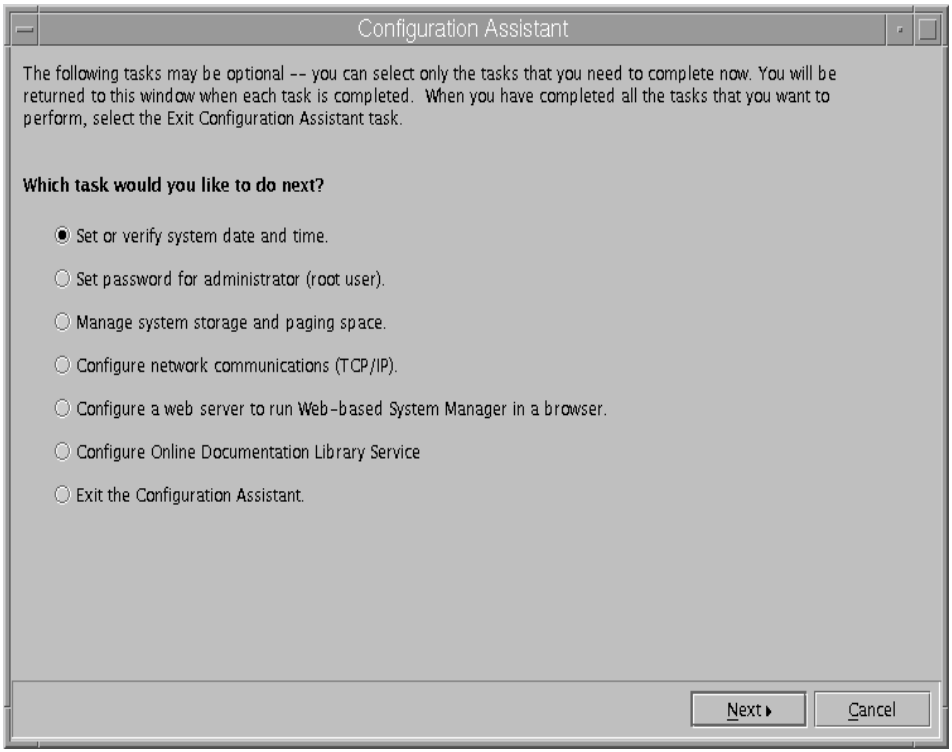

Figure 7-1 Configuration assistant menu

The snap command

In 4.5, “Documentation of existing environment” on page 57, we recommend using the **snap** command or a script to collect the most important system information. It is a good idea to run the **snap** command or your self-written script after the migration process to collect the same system information. This ensures that you can directly compare the output and identify any changes that might be of some concern to your production environment.

One example is that with AIX 5L some additional subsystems are active. This includes the Resource Monitoring and Control subsystem (RMC). Your task is it to compare the /etc/inetd.conf file with the original one running AIX Version 4.3. Hence, you need to decide whether or not these subsystems work in accordance with your security guidelines and if you want them to be active.

Furthermore, you can obtain this information when you compare the output of the **lssrc -a** command of your original AIX Version 4.3 system and your migrated AIX 5L system. Some of these activated daemons are needed if you want to benefit from new AIX 5L features, for example, System Hang Detection. In this

case, the associated daemons need to run in order to monitor the system at regular intervals. This means that obtaining a different output when comparing the system information of your original and your migrated system may not be a bad thing. However, we do recommend you verify each of these changes for yourself and decide whether or not they affect your production environment in any way.

Merging of AIX configuration files

After each migration installation, there is a final migration report that lists some configuration files that could not be merged. Some examples of these files (from one of our test migration scenarios) are shown in Example 7-1. It states that some configuration files could not be automatically merged into the system. The previous versions of these files are saved in the listed configuration directories.

You need to compare the saved files and the newly installed files to determine whether or not you need to recover configuration data. This means that you need to check the content of the configuration files and verify if there are new entries that you have to merge with your saved configuration files. If there are no new entries, you can just copy your backed up file over the new one. If you do need to merge the data, refer to the product documentation for a detailed description.

Example 7-1 Migration report

```
Configuration files which were saved in /lpp/save.config:
  /etc/3270.keys
  /etc/3270keys.hft
  /etc/bootptab
  /etc/inetd.conf
  /etc/mail/sendmail.cf
  /etc/map3270
  /etc/mib.defs
  /etc/ntp.conf
  /etc/rc.bsdnet
  /etc/rc.net
  /etc/rc.tcpip
  /etc/rpc
  /etc/services
  /etc/slip.hosts
  /etc/snmpd.conf
  /etc/snmpd.peers
  /etc/syslog.conf
  /etc/telnet.conf

 Configuration files which were saved in /usr/lpp/save.config:
  /usr/lpp/X11/bin/dynamic_ext
  /usr/lpp/X11/defaults/xinitrc
  /usr/lpp/X11/defaults/xserverrc
```

```
/usr/sbin/rsct/install/bin/ctposti
0513-071 The ctrmc Subsystem has been added.
0513-071 The ctcas Subsystem has been added.
```

Network configuration file

There are some important changes in the network and performance tuning parameters in AIX 5L Version 5.2. These include AIX kernel parameters and enhancements of the **no** and **nfso** command. Furthermore, you will find that the **vmtune** and **schedtune** commands are modified. These improvements ensure a more flexible and centralized system management. For a general discussion on enhanced network and performance parameters, refer to the *AIX 5L Version 5.2 Performance Management Guide*, which can be found at:

http://publib16.boulder.ibm.com/doc_link/en_US/a_doc_lib/aixbman/prftungd/prftungdtfrm.htm

We concentrate in this section on the issues involved when migrating your system to AIX 5L Version 5.2.

After you migrate your system from AIX Version 4.3 to AIX 5L Version 5.2, it is automatically set to run in compatibility mode. This means that the current behavior of the tuning commands is completely preserved. However, there is one exception, which is the **vmtune** parameter. This compatibility mode is controlled by a new *sys0* attribute that is called *pre520tune*. This will automatically be enabled as part of the migration process. To see the content of the current setting of the pre520tune attribute, you need to run the **lsattr -El sys0** command on your migrated AIX 5L Version 5.2 system.

When you disable the compatibility mode either using the Web-based System Manager, SMIT, or command line, you can see other changes. The examples shown in the list below are all related to the **no** command and include:

- arptab_bsiz
- arptab_nb
- extendednetstats
- ifsize
- inet_stack_size
- ipqmaxlen
- nstrpush
- pseintrstack

For more detailed information on the **no** command changes, refer to the *AIX 5L Version 5.2 Commands Reference*, which can be found at:

http://publib16.boulder.ibm.com/cgi-bin/ds_form?lang=en_US&viewset=AIX/

Trusted Computing Base (TCB)

If you have TCB enabled before the migration and want to check if it is still functioning properly after the migration, run the **tcpck - n** command. However, there is a problem issuing this command, which is a known problem that has workarounds. These error message is shown in Example 7-2.

Example 7-2 Enablement of TCB

```
3001-020 The file /dev/dlcqllc was not found
3001-089 The symbolic link from the file /usr/sbin/named to/usr/sbin/ named4
should not exist.
3001-089 The symbolic link from the file /usr/sbin/named-xfer to /usr/sbin/
named4-xfer should not exist.
```

To fix the problem, you need to run the commands shown in Example 7-3. The first line recreates the /dev/dlcqllc device. To correct the named4 errors you need to run the two commands shown in lines 2 and 3.

Example 7-3 Fixing of TCB problems

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

# tcbck -a /usr/sbin/named type=SYMLINK target=/usr/sbin/named4
# tcbck -a /usr/sbin/named -xfer type=SYMLINK target=/usr/sbin/named4-xfer
```

Security adjustments

Some security related files will be overwritten during the migration. This means that some of the security related adjustments in your environment might be lost after the migration process. We will describe some of the changes we encountered in our test environment. However, bear in mind that your production environment will have tighter security guidelines than our test environment. This means that there might be additional issues in addition to those mentioned below.

We therefore recommend that you check all security related files and adjustments as part of the post migration check list. This ensures that your system remains in accordance with the security guidelines of your company.

The following list contains changes we obtained in our test environment, which may affect your secured environment as well:

- The settings of /etc/security/user file are reset to the installation default. For example, the maxage parameter is set to 0, which means the password is valid without an end date. The other important value is maxretries, where the default value means you can have unlimited unsuccessful log in trials.

- The /etc/motd file is reset to the installation default value.
- For some security guidelines, you need an associated user and group name for every file. After the migration, the system file /usr/lpp/bos.net/inst_root/var/snapp has group ID177, which is not an allowed value.
- Entries in /etc/inetd.conf are different. In our test migration, the entries for ttdbserver and cmsd are reactivated after migration.
- No changes to snmpd.conf.
- After the migration to AIX 5l Version 5.2, ssh can still be used, but without using the PAM modules; you need to configure /etc/pam.conf as shown in Example 7-4, which is described on the following Web sites:

  http://www.ibm.com/servers/esdd/articles/openssh_updated.html
  http://publib16.boulder.ibm.com/pseries/en_US/aixbman/security/openssh.htm
  http://publib16.boulder.ibm.com/pseries/en_US/aixbman/security/securitytfrm.htm

Example 7-4 Example of /etc/pam.conf

```
shd auth required /usr/lib/security/pam_aix
OTHER auth required /usr/lib/security/pam_aix
sshd account required /usr/lib/security/pam_aix
OTHER account required /usr/lib/security/pam_aix
sshd password required /usr/lib/security/pam_aix
OTHER password required /usr/lib/security/pam_aix
sshd session required /usr/lib/security/pam_aix
OTHER session required /usr/lib/security/pam_aix
```

This means that the following modules are installed on the migrated machine:

- openssl-0.9.6b-1
- prngd-0.9.23-1
- openssh-2.9.9p2-5
- openssh-server-2.9.9p2-5
- openssh-clients-2.9.9p2-5

In comparison, on a newly installed AIX 5L Version 5.2 machine, the modules are:

- icacw005:root[/]> rpm -qa
- cdrecord-1.9-4
- mkisofs-1.13-4
- AIX-rpm-5.2.0.0-1
- prngd-0.9.23-2
- zlib-1.1.4-1
- openssl-0.9.6e-2
- openssl-devel-0.9.6e-2

– openssl-doc-0.9.6e-2

Mirrored disks

After the migration to AIX 5L, the bootlist of your AIX server will be set to the first disk. If you have a mirrored environment, the bootlist needs to be set to all disks belonging to the root volume group. This means that even though the boot record is already mirrored, you have to manually set the bootlist to both mirrored disks.

When migrating to AIX 5L, a new /dev/hd10opt file system mounted on /opt will automatically be created. This file system is created with the appropriate mirror copy, which means that there is no manual adjustment necessary. Another new file system is /proc. However, even though it is accessed like it is a file system, it is not physically on the disk, which means mirroring is not necessary.

Device driver removal

During the migration installation from AIX Version 4.3 to AIX 5L Version 5.2, all device drivers are removed. The reason for this is that some directories are removed during the migration process. These are listed below and include the following:

- /usr/lib/drivers
- /usr/lib/microcode
- /usr/lib/methods
- /dev

However, even though the new drivers for AIX 5L Version 5.2 are reinstalled during the migration process, the software support for device drivers must be reinstalled.

Non-root volume groups

As part of the post migration check, you need to ensure that all your user data from non-root volume groups is mounted correctly. In our test environment, we did not experience any problems importing these volume groups after the migration process. We recommend that you nonetheless check the order of entries in /etc/filesystems. This is an important issue, and you avoid having a mount point that is overmounted.

As an example, if our /data/images/dep1 file system is mounted first and then /data/images is mounted because of an entry order change in /etc/filesystems, you will not be able to access any data from /data/images/dep1. If you, however, change the order to the correct one again and mount /data/images first and then /data/images/dep1, you can access all system data, as was possible before the migration process.

This means that before you import your user volume groups, check the sequence in /etc/filesystems and adjust it before mounting your data file systems.

7.2 NIM master specific tasks

For the majority of AIX servers, you will be able to use any of the migration methods introduced in Chapter 5, “Migration by media” on page 79 and Chapter 6, “Migration by NIM” on page 105. For the NIM master, however, the only way to do the migration is to use media. In this case, the configuration of the NIM master remains unchanged. The important task is then to create AIX 5L NIM resources on this machine. This is explained in detail in 6.3, “NIM master migration” on page 117.

Alternatively, you can migrate your NIM master and move it to a different hardware that supports AIX 5L Version 5.2. In this case, you need to perform some reconfigurations on the NIM master machine. Refer to 6.3.4, “Moving the NIM master from POWER3 to POWER4” on page 127.

In this section, however, we concentrate on post migration tasks that are specific to the NIM master only. We recommend that you verify all discussed issues of the common post migration section of Chapter 7, “Post migration tasks” on page 165 as well.

NIM master filesets

As part of the AIX migration process, the obsolete filesets will be removed from your system and not be replaced. This includes some of the AIX Version 4.3 NIM master filesets, which we list in Table 7-1.

Table 7-1 Removed NIM master filesets during AIX migration

Fileset name	Description
bos.msg.en_US.sysmgt.nim.master_gui	NIM GUI Messages - US
bos.sysmgt.nim.master_gui	Network Install Manager - GUI

Some of these filesets will be replaced when performing the migration installation on your NIM master. Table 7-2 on page 174 lists these replaced AIX 5L Version 5.2 NIM master filesets.

Table 7-2 NIM master filesets replaced after AIX migration

Fileset name	Description
X11.Dt.helpnim	AIX CDE Minimum Help File
X11.msg.en_US.Dt.helpnim	Network Install Manager
bos.sysmgt.nim.client	Network Install Manager
bos.sysmgt.nim.master	Network Install Manager
bos.sysmgt.nim.spot	Network Install Manager - SPOT

The NIM master configuration

Even though we migrate the NIM master to AIX 5L Version 5.2 we can still use this machine as NIM master AIX Version 4.3 as before. To verify that the original AIX Version 4.3 NIM environment is still configured you need to check the /etc/niminfo file. One important requirement is that the network settings have not changed which means the NIM master must be able to communicate with all client machines. In Example 7-5 we show a sample output of /etc/niminfo taken from the migrated NIM master of our test environment.

Example 7-5 Verification of /etc/niminfo

```
# pg /etc/niminfo

export NIM_NAME=master
export NIM_CONFIGURATION=master
export NIM_MASTER_PORT=1058
export NIM_REGISTRATION_PORT=1059
export NIM_MASTER_HOSTNAME=srvr50f.itsc.austin.ibm.com
```

For a detailed output of the NIM environment configuration, use the **lsnim -l** command. Additionally, you need to check the NIM resources using the following command:

```
# nim -o check <resource name>
```

Rebuilding the /etc/niminfo file

Even though we did not encounter any problems with the /etc/niminfo file after the migration to AIX 5L Version 5.2, we give you some information on how to rebuild it. This will just be of any interest to you if you do encounter a corrupted /etc/niminfo file. For the rebuild, you obtain the information stored in the NIM database, which is used to perform NIM commands and operations. You can choose between performing the rebuild using the Web-based System Manager, SMIT, or command line, as shown in the following sections.

Option 1: Using the Web-based System Manager

To rebuild the /etc/niminfo using the Web-based System Manager, type **wsm** on the command line and follow the steps below and you will see the window shown in Figure 7-2.

1. Select **Network Installation Management**.
2. Select **NIM** -> **Advanced Configuration** -> **Rebuild Master Configuration File**.

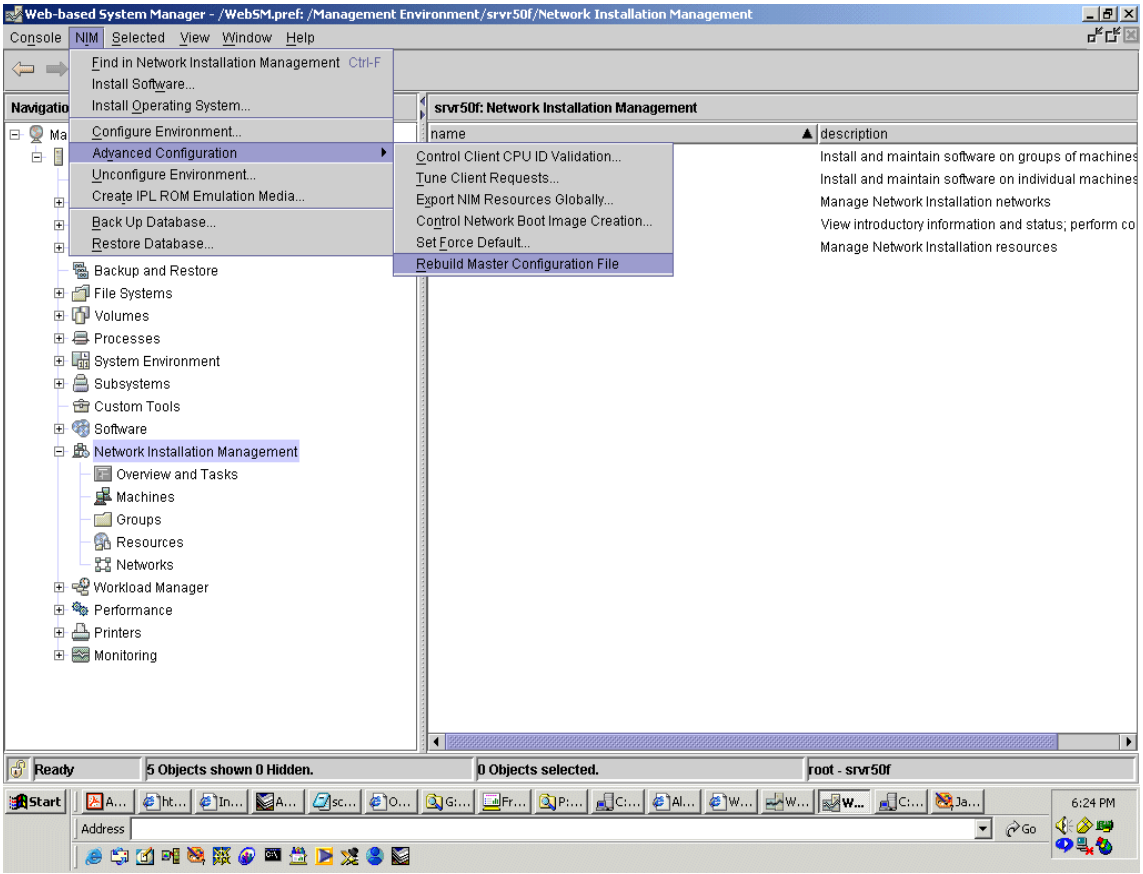

Figure 7-2 Rebuilding /etc/niminfo file using the Web-based System Manager

Option 2: Using SMIT

Using SMIT to rebuild the /etc/niminfo file on the NIM master, follow the outlined steps below. You should get the menu shown in Example 7-6 on page 176. Type **smitty** on the command line and select **Software Installation and Maintenance** -> **Network Installation Management** -> **Perform NIM Administration Tasks**.

Example 7-6 Rebuilding the /etc/niminfo file

```
                  Perform NIM Administration Tasks

Move cursor to desired item and press Enter.

Manage Networks
Manage Machines
Manage Resources
Manage Groups
Backup/Restore the NIM Database
Configure NIM Environment Options
Rebuild the niminfo File on the Master
Unconfigure NIM



F1=Help             F2=Refresh       F3=Cancel        F8=Image
F9=Shell            F0=Exit          Enter=Do
```

Option 3: Using command line

From the command line, use the **nimconfig** command with the -r flag as shown below:

```
# nimconfig -r
```

7.3 Alternate disk migration specific tasks

In this section, we describe specific post migration tasks that are necessary if you migrate your AIX system using the alternate disk migration method. In this case, the migration takes place on a spare disk on your system. This volume group after completing the migration process is named altinst_rootvg. Your original disk is still called the root volume group and is activated by running AIX Version 4.3.

To swap from your AIX Version 4.3 rootvg to your cloned and migrated altinst_rootvg with AIX 5L as the installed operating system, you need to stop all your applications, change the bootlist to the AIX 5L disk, and initiate the reboot of your machine. You then need to go through the list of common post migration tasks, as we describe at the beginning of this chapter. After completing these tasks, your AIX 5L system is ready to be part of your production environment again.

However, you might want to put your migrated rootvg back on hdisk0. This is not necessary but it means that you can use the spare disk for other purposes. You

first need to remove the original AIX Version 4.3 disk, which is now called old_rootvg. For this task, use the command below:

```
# alt_disk_install -x old_rootvg
```

This is the recommended method, even though you could alternatively use the **exportvg** command. However, we advise you to test this command in advance in your environment and check if the content of your /etc/filesystems is still consistent after running the **exportvg** command.

Once you have removed the old_rootvg on hdisk0, you can mirror the active rootvg to hdisk0, change the bootlist, and things are as they would be if you migrated your rootvg directly without the cloning process.

8

Applications

When you migrate your operating system from AIX Version 4.3 to AIX 5L, you need to verify whether or not your applications will run with the new version of the operating system as before. You need to verify if it is necessary to upgrade the application release at the same time as the operating system migration. In other cases, you may need to only apply patches.

Another issue is moving your application to a different hardware platform as well. This may include transferring your data with external disks by first exporting the data and then importing it onto the new hardware running AIX 5L. One of the most important tasks for you to do is to refer to your applications vendor documentation. This way, you will find detailed upgrade procedures that will help you migrate your applications safely and with minimized downtime from AIX Version 4.3 to AIX 5L. Most importantly, allow enough time for porting your applications to AIX 5L and consider setting up test environments prior to starting the migration of your production environment. This is going to be the most crucial part of your migration process.

In the next chapters, we concentrate on some of the most used applications at customer sites. We will not describe a detailed upgrade procedure for each application, but we want to give some guidelines of the order of the migration process, things to be aware of and to look out for, and necessary tasks to do to prepare your system.

8.1 SAP and Oracle

In this section, we will not describe a detailed upgrade of SAP or Oracle, but concentrate on the AIX related issues involved when migrating your AIX Version 4.3 system running Oracle 8.1.7 to AIX 5L Version 5.1 running Oracle 9.2. The information we provide is complementary to the documentation of the official SAP Guide *Upgrade to Oracle Version 9.2.0: UNIX*. Furthermore, we will not cover every issue involved for each customer situation. This information acts as a guideline of things to consider when planning your migration process. We discuss how and in what order to do the AIX migration and Oracle upgrade and outline post-migration issues. In particular, customers running SAP/Oracle in 64-bit mode need to carefully plan and prepare the migration to AIX 5L. The reason for that is that some necessary operations, such as increasing the system table space, cannot be done after the operating system migration is completed.

Another important issue for most customers is the minimization of the system's downtime. There are several pre-migration steps you can perform on the original AIX Version 4.3 system in order to keep the downtime low when migrating your operating system to AIX 5L and your database to Oracle 9.2. We will discuss certain issues involved in moving hardware platforms. This involves moving from a POWER3 machine running AIX Version 4.3 to a POWER4 machine, like IBM @server pSeries p630, p650, p655, p670, or p690 running AIX 5L as well as SAP and Oracle.

Another good source of information related to SAP and AIX is the following publication, *A Holistic Approach to a Reliable Infrastructure for SAP R/3 on AIX*, SG24-5050. This includes AIX shared memory management and other performance related issues.

At the time of the writing of this redbook, the 64-bit product availability for AIX, DB2 UDB, and Oracle DB is as shown in Table 8-1 on page 181. SAP has certified AIX 5L Version 5.1 and 5.2 of the AIX 5L family. For further information, refer to SAP Note 502532. For information on how to access SAP Notes, refer to 8.1.1, "Pre-migration issues" on page 181. Detailed information about which SAP component is available for each of the database and Web Application Server versions can be obtained from the infoservice at the IBM SAP International Competence Center (ISICC). You need to contact isicc@de.ibm.com.

Table 8-1 SAP DB and APP server on pSeries 64-bit availability

Operating system	Database	SAP Basis/Web Application Server
AIX Version 4.3	Oracle 8.1.7	4.6D, 6.1, 6.2
	DB2 UDB V7	4.6D, 6.1, 6.2
AIX 5L Version 5.1 and 5.2	Oracle 9.2	Oracle 9.2
	Oracle RAC	6.2
	DB2 UDB V7	4.6D, 6.2
	DB2 UDB V8	6.2

8.1.1 Pre-migration issues

Before you start planning and preparing the migration process, read this section carefully and refer to the SAP notes. These can be found at the SAP Service Marketplace at the following Web sites:

http://service.sap.com/instguides
http://service.sap.com/notes
http://service.sap.com/systemmanagement

To access the Web sites of SAP Service Marketplace, you need an SAPNet user ID (which can be obtained at http://service.sap.com). Some of the relevant SAP Note numbers are 540 021, 539 921, 539 970, and 539 922. Furthermore, this is the place where you find the official Oracle 9.2.0 Migration guide called *Upgrade to Oracle Version 9.2.0: UNIX*.

As with the AIX migration, when you plan and prepare your migration process, ensure that you have all the needed software ordered and available. This includes AIX 5L and appropriate maintenance levels as well as the SAP Update Kits:

- OS Update Kit AIX 5L Version 5.1 64-bit 4.6D Kernel
- OS Update Kit AIX 5L Version 5.1 64-bit 6.2.0 Kernel

Additionally, you need to download Oracle patches and patch sets, as described in the guide book *Upgrade to Oracle Version 9.2.0: UNIX*.

32-bit versus 64-bit

It is possible to execute both 32-bit and 64-bit software on the 32-bit as well as on the 64-bit AIX kernel. The advantage of the 32-bit AIX kernel is the ability to deal with existing resources. The 64-bit AIX kernel allows you, however, to

access more than 96 GB of RAM. For further discussions on this topic, refer to 4.2.2, “Binary compatibility” on page 43.

8.1.2 Migration and upgrade procedure

In this section, we describe the steps necessary to migrate your system from AIX Version 4.3 to AIX 5L Version 5.1, as well as outline the Oracle upgrade procedure when moving your database from Version 8.1.7 to 9.2.0. We begin with some preparation tasks that you should do on your AIX Version 4.3 system prior to the migration process. In the next step, we prepare the SAP kernel, do the Oracle upgrade, and start the AIX migration itself.

AIX prerequisites

The prerequisites for AIX 5L Version 5.1 and Version 5.2 are listed in the following sections.

AIX 5L Version 5.1

For the post-migration, you will need the latest maintenance level of AIX. You can reduce the total downtime of your system if you already download it at this stage. The latest maintenance level for AIX 5L Version 5.1 is (at the time of writing) ML 03 which you can download at the following URL by searching for APAR IY32749:

`http://techsupport.services.ibm.com/server/aix.fdc`

Additionally, you need to check and see if you have the following APARs installed and download them using the same URL if needed:

- IY28949: Array pointer incorrectly used in aio_nwait
- IY28766: Add bread/bwrite, lread/lwrite, and phread/pwrite for
- IY30150: waitpid dumps core on 64-bit, when called from escepti
- IY29965: THREAD_WAIT() computes wait time incorrectly
- IY26778: truss does not report the coded Error Numbers

The way to check if they are installed on your system is by using the **`instfix`** command, as shown in Example 8-1 on page 183. The first two lines represent the case where you do not have to install the command because it already is on your machine. The second case demonstrates the output you get when it is not on your system. In this case, you need to download the fix.

Example 8-1 AIX fixes installation check

```
# instfix -ik IY28949
  All filesets for IY28949 were found.

# instfix -ik IY28949
  There was no data for IY28949 in the fix database.
```

AIX 5L Version 5.2

There is no official guidelines for AIX 5L Version 5.2, but the prerequisites should be the same and the upgrade works the same way.

Additionally, you have to fulfil some disk space requirements. These include 4 GB in $ORACLE_HOME, 3.6 GB in /oracle/stage, 400 MB in /tmp, and 160 MB in /oracle. If you plan to transfer your data using external disks, you can increase the file system size as required before starting the migration process from AIX Version 4.3 to AIX 5L.

After you have migrated your system from AIX Version 4.3 to AIX 5L, you will not be able to start and use Oracle 8.1.7 64-bit anymore. Therefore, it is important to follow the order of preparation steps that we describe in the following paragraphs. This means that you should first prepare the SAP kernel and then do the database upgrade, follow the preparation steps on AIX Version 4.3 and then do the migration of the operating system. This ensures a successful migration with as little downtime as possible. This applies to the 64-bit version of Oracle only and is not an issue with the 32-bit version. For a detailed discussion, refer to “Oracle upgrade” on page 184.

SAP kernel preparation

As mentioned before, when discussing the minimization of the downtime, there are several steps you can do and prepare on your running AIX Version 4.3 system. These include preparing the SAP kernel 4.6D as well as kernel 6.20 for AIX 5L Version 5.1. However, this kernel can only be prepared but not used on AIX Version 4.3. We recommend that you install the new kernel in /sapmnt/<SAPSID>/exe_920_51, for which you need 300 MB of disk space. You can then prepare the kernel in this directory and, after the AIX migration, all you have to do is switch the old and new kernel by renaming the directories.

As one of the pre-migration tasks, you download the SAP package OS Update Kit AIX5.1 64-bit 4.6D Kernel as well as OS Update Kit AIX 5L Version 5.1 64-bit 6.2.0 Kernel. This package contains an SAP kernel CD and all required components needed to set up the AIX 5L enabled 64-bit SAP kernel.

An outline of steps to follow in order to prepare the SAP kernel on your running AIX Version 4.3 system is as follows:

1. Log on as user <SAPSID>.
2. Mount the SAP kernel CD.
3. Unpack the kernel files into /sapmnt/<SAPSID>/exe_920_51.
4. Extract the SAP DBATOOLS (do not use DBATOOLS from the SAP kernel CD).
5. Change the permissions of the DBA tools.
6. Change the authorization for the programs sapdba, brarchive, brbackup, brconnect, and brtools.
7. Change the permission of saposcol to 4755.
8. Upgrade the startdb/stopdb scripts.

For detailed instructions, refer to the SAP documentation.

Completing all the tasks ensures that the SAP kernel is prepared. This way, we can easily switch it after the operating system migration from AIX Version 4.3 to AIX 5L and the database upgrade from Oracle 8.1.7 to 9.2.0.

Oracle upgrade

There are different ways to perform the Oracle Upgrade from Version 8.1.7 to 9.2.0. If you are running a 32-bit SAP/Oracle 8.1.7 environment, you can migrate your server to AIX 5L Version 5.1 and keep running the 32-bit SAP and Oracle as before. You can perform the database upgrade to Oracle 9.2.0 at a later time.

If you are running 64-bit SAP/Oracle 8.1.7 on your AIX Version 4.3 system, you are required to do the database upgrade to Oracle 9.2.0 and the operating system migration to AIX 5L Version 5.1 at the same time. The reason for this is that the 64-bit executable of AIX Version 4.3 will not run on AIX 5L without recompilation, as described in 4.2.2, “Binary compatibility” on page 43.

If you do not have strict requirements on the downtime of your AIX server, follow the SAP Database Upgrade guide *Upgrade to Oracle Version 9.2.0: UNIX*. Additionally, you will find comments, remarks, and suggestions on some of the steps described in IBM SAP Technical Brief *Migration of Oracle 8.1.7/AIX 4.3.3 SAP installation To Oracle 9.2/AIX 5.1 SAP Installation* (Version 1.00, December 2002). This can be obtained by writing to isicc@de.ibm.com. Some of the steps described have been done already when you prepared the SAP kernel. This includes upgrading the SAP Database Administration Tools and startdb/stopdb scripts.

Another option is to change the hardware at the same time as doing the operating system migration and the database upgrade. This reduces the total amount of your systems downtime again. For this task, you need a new machine running AIX 5L Version 5.1 in addition to your original AIX Version 4.3 system. You can do the Oracle installation process independently from all other steps. An outline of the order in which you have to do this upgrade and migration process is as follows:

1. Perform a new and complete installation of AIX 5L on your new hardware.
2. Install AIX maintenance levels and fixes as described before.
3. Set up the network. For this task, we recommend using the same host name as before.
4. Recreate the SAP/Oracle users and groups manually on the AIX 5L machine.
5. Transfer all files in /home/<sapsid>adm to the new machine.
6. Add all SAP related entries to /etc/services on the new machine.
7. Shut down Oracle and SAP on your old machine.
8. Back up the database
9. Transfer all files from /oracle, /sapmnt, and /usr/sap and below to the new machine:
   - Use external disks: Unmount file systems of the volume group, export and varyoff the volume group, disconnect disks from the old machine, reconnect to the new machine, and import the volume group. The file transfer is complete.
   - Manual recreation of the file systems: Transfer data via the network or tape from the old to the new machine.
10. Do the database upgrade to Oracle 9.2.0
11. Perform all necessary post-migration steps as before.

Preparation of the AIX system

Important: After the migration from AIX Version 4.3 to AIX 5L, you will not be able to start and use Oracle 8.1.7 64-bit anymore.

You need to do some preparation steps on AIX Version 4.3 after the upgrade of Oracle to 9.2.0 and before starting the operating system migration.

The first step involves the kernel extension. Older releases of Oracle include kernel extensions (pw-syscall) that are loaded at startup time. This worked through an entry in the /etc/inittab. Oracle 9.2.0 on AIX 5L does not need these kernel extensions anymore, but they are not removed automatically by the

Oracle upgrade process. That means that you need to remove the following entry manually from the /etc/inittab:

```
orapw:2:wait:/etc/loadext -l /etc/pw-syscall
```

In addition to all the steps performed in the Oracle upgrade process, do the following:

- Check the database requirements (CHECKS.SQL).
- Extend the System Tablespace.
- Check for Invalid Database Objects.
- Run the PREMIG.SQL script.

These steps are described in detail in Chapter 3 of the SAP Database Upgrade Guide *Upgrade to Oracle Version 9.2.0: UNIX*.

You can now start your AIX migration using any of the methods described in Chapter 5, “Migration by media” on page 79 and Chapter 6, “Migration by NIM” on page 105 of this book. This means that you can either migrate your AIX system by media, using the Network Installation Management (NIM) features, or by using an alternate disk migration.

8.1.3 Post-migration issues

After the preparation of the SAP kernel, the upgrade of the database to Oracle 9.2.0, and the migration of the operating system to AIX 5L, you need to complete your migration process by performing some additional steps. The first one involves installing the latest maintenance level for AIX 5L. For AIX 5L Version 5.1, it is ML 03, as described in “AIX prerequisites” on page 182. For AIX 5L Version 5.2, there is no maintenance level available at the time of writing. Furthermore, some filesets are required by Oracle 9i and SAP. This means that you need to install the base levels and download and install the following required filesets:

- bos.adt
- bos.perf.libperfstat
- Java.rte 1.3.1
- perl.rte (you need Perl 5.6.1 or later)

The next step is to switch the SAP kernel and active the new one that you prepared in “SAP kernel preparation” on page 183. We recommend you keep the old kernel in case you have application instances still running on AIX Version 4.3 servers. You can share your SAP kernel via NFS, which means you might be able to continue using the AIX Version 4.3 code. To activate the new AIX 5L Version 5.1 kernel, you need to follow these steps:

1. Log on as <sapsid>adm.
2. Change to the /sapmnt/<SAPSID> directory.

3. Rename the old AIX Version 4.3 kernel directory.
4. Activate the new AIX 5L Version 5.1 kernel.

The next step is to reconfigure the Oracle Client Software, as described in the guide *Upgrade to Oracle Version 9.2.0: UNIX* and then you can restart R/3. There are some additional updating and checking parameter steps necessary on the database, but this can be done while the system is up and running. The final step is to back up the database.

8.2 DB2

Our second application section deals with the topic of DB2 UDB. This issue is of concern to you when you have version of DB2 running on your AIX Version 4.3 system that you want to migrate to AIX 5L. The most important information you need to obtain is the currently supported release versions of DB2 on AIX Version 4.3 and AIX 5L. The latest versions, at the time of the writing of this redbook, are summarized in Table 8-2. Additionally, you need to verify which DB2 and AIX kernel versions are supported. This is listed and explained Table 8-2 as well.

Table 8-2 DB2 supported on AIX release matrix

DB2	AIX Version 4.3	AIX 5L Version 5.1	AIX 5L Version 5.2
V7	YES	YES	YES
V8	NO	YES	YES

DB2 Version 7 comes with three CDs, which include the following:

- Filesets 7.1.0.xx: 32-bit DB2 instances on AIX Version 4.3 and AIX 5L 32-bit kernel
- Filesets 7.1.1.xx: 64-bit and/or 32-bit DB2 instances on AIX Version 4.3 and AIX 5L 32-bit kernel
- Filesets 7.1.2.xx: 64-bit and/or 32-bit DB2 instances on AIX 5L 64-bit kernel

For DB2 Version 7 and Version 8, you need to check the product release notes to determine the required minimum fixpack levels for the product. We recommended you obtain the latest service. Both DB2 and AIX will require service to be installed over the base product levels.

DB2 V8 requires AIX 5L Version 5.1 ML02, including the following APARs: IY31254, IY32217, IY32905, IY33023, and IY29345. If you are using AIX 5L Version 5.2, these APARs are included automatically. DB2 V8 (8.1.0.8) is for the 32-bit AIX Version 4.3 system. DB2 V8 (8.1.1.8) can be used for the 32-bit and 64-bit AIX 5L installation.

Furthermore, in Table 8-3, we give you the end of support dates for the various versions of DB2 UDB. This means the data of the DB2 database and DB2 connect defect support expires. You can in some cases extend the support for a additional fee.

Table 8-3 DB2 service matrix

Version	Edition	End of service
7.1	EE, EEE, and Connect EE to NUMAQ	06/30/2003
7.2	All editions	03/31/2004
8.1	All editions	03/31/2005

For more information on DB2 related issues, refer to the following URL:

http://www.ibm.com/cgi-bin/db2www/data/db2/udb/winos2unix/support/newsletter.d2w/n20020101

Test environment setup

In this section, we want to share our experience of migrating an AIX Version 4.3 system running DB2 Version 7.1 with the SAMPLE database and default instances. Our test machine is an RS/6000 Model F80. In Example 8-2, we show an output of the DB2 filesets, including their version number, after running the **lslpp** command.

Example 8-2 DB2 filesets

```
  db2_07_01.cdb            7.1.1.2  COMMITTED  Control Database
  db2_07_01.cj             7.1.1.2  COMMITTED  Java Common files
  db2_07_01.client         7.1.1.2  COMMITTED  Client Application Enabler
  db2_07_01.cnvucs         7.1.1.2  COMMITTED  Code Page Conversion Tables -
  db2_07_01.conn           7.1.1.2  COMMITTED  Connect
  db2_07_01.conv.jp        7.1.1.2  COMMITTED  Code Page Conversion Tables -
  db2_07_01.conv.kr        7.1.1.2  COMMITTED  Code Page Conversion Tables -
  db2_07_01.conv.sch       7.1.1.2  COMMITTED  Code Page Conversion Tables -
  db2_07_01.conv.tch       7.1.1.2  COMMITTED  Code Page Conversion Tables -
  db2_07_01.cs.drda        7.1.1.2  COMMITTED  Communication Support - DRDA
  db2_07_01.cs.ipx         7.1.1.2  COMMITTED  Communication Support - IPX
  db2_07_01.cs.rte         7.1.1.2  COMMITTED  Communication Support -
  TCP/IP
  db2_07_01.cs.sna         7.1.1.2  COMMITTED  Communication Support - SNA
  db2_07_01.das            7.1.1.2  COMMITTED  Administration Server
  db2_07_01.db2.engn       7.1.1.2  COMMITTED  Engine
  db2_07_01.db2.rte        7.1.1.2  COMMITTED  Run-time Environment
  db2_07_01.db2.samples    7.1.1.2  COMMITTED  Sample Database Source
  db2_07_01.dj             7.1.1.2  COMMITTED  Distributed Join for DB2 Data
```

```
  db2_07_01.elic           7.1.1.2  COMMITTED  Product Signature for UDB
  db2_07_01.gs             7.1.1.2  COMMITTED  Getting Started
  db2_07_01.jdbc           7.1.1.2  COMMITTED  Java Support
  db2_07_01.ldap           7.1.1.2  COMMITTED  DB2 LDAP Support
  db2_07_01.msg.En_US      7.1.1.2  COMMITTED  Product Messages - U.S.
  db2_07_01.repl           7.1.1.2  COMMITTED  Replication
  db2_07_01.tspf           7.1.1.2  COMMITTED  Transformer Stored Procedure
  db2_07_01.wcc            7.1.1.2  COMMITTED  Control Center
```

We describe in Chapter 5, “Migration by media” on page 79 and Chapter 6, “Migration by NIM” on page 105 various ways to migrate your AIX Version 4.3 system. The most straight forward method is using media. The most complex method is using alternate disk migration. However, this method also guarantees the shortest downtime and therefore least interruption for our database environment. We therefore decided to use the alternate disk migration to move our AIX Version 4.3 system to AIX 5L Version 5.2. We used one of the NIM masters, which was already set up in our environment for some of the earlier test scenarios. For further information on this migration method, refer to 6.4.3, “Alternate disk migration of client” on page 140.

For our test environment, the use of the alternate disk migration means that during the migration process, the database is still running as normal. We show a list of active databases in Example 8-3. This command lists the database name and path we use in our test environment.

Example 8-3 Active databases

```
db2 => list active databases

                            Active Databases

Database name                              = SAMPLE
Applications connected currently           = 1
Database path                              = /home/db2inst1/db2inst1/NODE0000
/SQL00001/
```

Some DB2 applications require a certain level of Java filesets. With AIX 5L, Java Version 1.3.1 will be installed, as shown in Example 8-4 on page 189. Before starting the migration process, you can remove the old Java filesets. That is what we did, so you only can only see the Java Version 1.3.1 filesets.

Example 8-4 Java Runtime Environment on AIX 5L Version 5.2

```
# lslpp -l | grep Java

Java131.rte.bin            1.3.1.2  COMMITTED  Java Runtime Environment
Java131.rte.lib            1.3.1.2  COMMITTED  Java Runtime Environment
```

```
db2_07_01.cj              7.1.1.2  COMMITTED  Java Common files
db2_07_01.jdbc            7.1.1.2  COMMITTED  Java Support
```

As mentioned before, one important advantage of the alternate disk migration is that your application can run during this process. This is shown in Example 8-5, with a sample output of the DB2 processes that were active during the migration process.

Example 8-5 DB2 processes during AIX migration

```
# ps -ef | grep -i db2

db2inst1  6090  4266   0 10:47:46  pts/0  0:00 -ksh
    root 13010     1   0 10:37:29      -  0:00 db2wdog
db2inst1 13960 15998   0 10:37:27      -  0:00 db2tcpcm
db2inst1 14300 15454   0 10:49:43      -  0:00 db2agent (SAMPLE)
db2inst1 14374  6090   0 10:47:48  pts/0  0:00 db2
db2inst1 14990 16382   0 10:37:28      -  0:00 db2resyn
db2inst1 15234 15998   0 10:37:27      -  0:00 db2tcpcm
db2inst1 15454 15998   0 10:37:27      -  0:00 db2ipccm
    root 15810     1   0 10:37:27      -  0:00 db2wdog
db2inst1 15998 15810   0 10:37:27      -  0:00 db2sysc
db2inst1 16382 15998   0 10:37:27      -  0:00 db2gds
db2inst1 16534 16382   0 10:49:43      -  0:00 db2dlock
db2inst1 16808 16382   0 10:49:43      -  0:00 db2pfchr
db2inst1 17032 16382   0 10:37:28      -  0:00 db2srvlst
db2inst1 17290 16382   0 10:37:28      -  0:00 db2spmlw
db2inst1 17548 15998   0 10:37:28      -  0:00 db2spmrm
db2inst1 17908 15454   0 10:49:35      -  0:00 db2agent (instance)
db2inst1 18086 16382   0 10:49:43      -  0:00 db2loggr
   db2as 18838 13010   0 10:37:29      -  0:00 db2sysc
   db2as 19092 18838   0 10:37:29      -  0:00 db2gds
   db2as 19350 18838   0 10:37:29      -  0:00 db2ipccm
   db2as 19608 19092   0 10:37:29      -  0:00 Scheduler
   db2as 19866 18838   0 10:37:29      -  0:00 db2tcpcm
   db2as 20124 18838   0 10:37:29      -  0:00 db2tcpcm
   db2as 20382 18838   0 10:37:29      -  0:00 db2tcpdm
db2inst1 20640 16382   0 10:49:43      -  0:00 db2pfchr
db2inst1 20898 16382   0 10:49:43      -  0:00 db2pfchr
db2inst1 21156 16382   0 10:49:43      -  0:00 db2pclnr
    root 22538 13484   0 13:52:19  pts/1  0:00 grep -i db2
```

After the alternate disk process is completed, we stop the database and reboot the AIX machine for the cloned and migrated disk. We complete the post migration tasks as described in Chapter 7, “Post migration tasks” on page 165 and started DB2 Version 7 again. The database ran without problems, as we expected. The advantage is that we used a database version that is supported

on AIX Version 4.3 and on AIX 5L. If you want to change database fileset versions at the same time, your tasks would be slightly different. This task requires you to deinstall the database first and then to reinstall the new filesets after the AIX migration.

8.3 HACMP

The High Availability Cluster Multiprocessing (HACMP) software provides numerous facilities you can use to build highly available clusters. When migrating from AIX Version 4.3 to AIX 5L, you need to understand which version of HACMP is supported at each release level of AIX. Table 8-4 gives you an overview of HACMP support for AIX Version 4 and 5.

Table 8-4 Support matrix for HACMP

		AIX operating system version		
		4.3.3	**5.1**	**5.2**
HACMP Classic and ES	4.4.0	Y	Y	N
	4.4.1	Y	Y	-
	4.5	N	Y	Y

Note that HACMP Version 4.4.0 has been withdrawn from service. Furthermore, at the moment, only the Enhanced Scalability (ES) Version of HACMP is supported in LPARs if you run your server in partition mode.

HACMP Version 4.5 requires AIX 5L Version 5.1 and a minimum level of RSCT Version 2.2.1.0. We also recommend you have your AIX 5L maintenance level up to date. Furthermore, the filesets listed below are required to successfully use HACMP:

- bos.adt.libm
- bos.adt.syscalls
- bos.data
- rsct.compat.basic.hacmp
- rsct.compat.clients.hacmp
- vacpp.msg.en_US.ioc.rte
- vacpp.ioc.aix50.rte
- csm.client
- devices.chrp.base.ServiceRM

In addition, you are required to have the following APARs applied to guarantee system support:

- HACMP HAS and CRM:
  - HA 4.5 APAR IY36938
  - AIX 5L Version 5.2 APAR IY36782, IY37744, and IY37746
- HACMP ES and ESCRM:
  - HA 4.5 APAR IY36938, IY36933
  - AIX 5L Version 5.2 APAR IY36782, IY37744, IY37746, and IY36626

8.4 IBM Tivoli® Storage Manager (TSM)

IBM Tivoli Storage Manager (TSM) protects your system's data from unexpected hardware failures and errors by storing backup and archive copies of data on offline storage. In this section, we will not describe detailed upgrade steps of TSM, but introduce AIX related issues involved when migrating your AIX Version 4.3 system to AIX 5L.

IBM Tivoli Storage Manager provides three types of AIX installation packaging. The type used depends on the content of the delivery. The following list introduces these installation methods:

- New install and migration

  For a new version or release of Tivoli Storage Manager (TSM), install packages are set up to do a new installation if no prior version is installed on the system. Alternatively, you can migrate from the earlier version.

- Update install

  An update install package is used with a maintenance update (PTF) or a patch is provided to correct software problems.

If you have TSM Version 4.2 installed on your AIX Version 4.3 system and you migrate to AIX 5L Version 5.1, there are device driver conflicts. To solve this problem, you need to install the Tivoli Storage Manager device support for AIX 5L Version 5.1, as listed below:

```
tivoli.tsm.devices.aix5.rte
```

For detailed and updated information on IBM Tivoli Storage Manager, refer to the following URL:

http://publib.boulder.ibm.com/tividd/td/IBMStorageManagerforAIX5.1.html

TSM Version 5.1.5 (server and client) and higher is supported with AIX Version 4.3 and AIX 5L Version 5.1 and 5.2. Table 8-5 and Table 8-6 on page 193 give an overview of the supported versions of the IBM Tivoli Storage Manager Version 4.2 and 5.1 regarding AIX server requirements and AIX client requirements.

Table 8-5 AIX Server minimum hardware and software requirements

Hardware	Operating system	Communication protocol
Version 5.1 RS/6000, either 32- or 64-bit or pSeries with at least 128 MB of RAM and 120 MB of disk space	Version 5.1 IBM AIX Version 4.3.3 (32-bit) or later, or IBM AIX 5L Version 5.1 (32-bit or 64-bit) or later	Version 5.1 ► TCP/IP, which comes standard with AIX ► Shared Memory Protocol (for AIX Client only)
Version 4.2 Minimum hardware required for AIX Version 4.3.3 or IBM AIX 5L Version 5.1 or later with at least 128 MB of RAM and 120 MB of disk space	Version 4.2 AIX Version 4.3.3 or IBM AIX 5L Version 5.1 or later	Version 4.2 ► TCP/IP, which comes standard with AIX ► Shared Memory Protocol (for AIX client only)

Table 8-6 AIX client minimum hardware and software requirements

Hardware	Operating system	Communication protocol	Additional software
Version 5.1 ► RS/6000 or pSeries workstation with at least 128 MB of RAM and 120 MB of disk space ► The HSM client requires an additional 256 MB of RAM (512 MB recommended) and 1GB of disk space for an HSM managed file system (10 GB is recommended)	Version 5.1 One of the following: ► AIX Version 4.3.3 (32-bit) ► AIX 5L Version 5.1 (32-bit or 64-bit) ► AIX 5L Version 5.2 (32-bit or 64-bit) For the HSM client: ► AIX Version 4.3.3 (32-bit) ► AIX 5L Version 5.1 (32-bit or 64-bit)	Version 5.1 TCP/IP, which comes standard with AIX Shared Memory Protocol (with AIX server only)	Version 5.1 The backup-archive client requires: ► X Windows System, X11R6 ► Motif 1.2 or 2.0 ► CDE ► JFS The HSM client requires: ► X Windows System, X11R6 ► Motif 1.2 or 2.0 ► CDE ► JFS

Hardware	Operating system	Communication protocol	Additional software
Version 4.2 RS/6000 or pSeries workstation with at least 64 MB of RAM 8 MB of available disk space, and 3 MB of additional disk space for the API client The HSM client needs at least 256 MB of RAM (512 MB recommended) and 1GB of disk space for an HSM managed file system (10 GB is recommended)	**Version 4.2** AIX Version 4.3.3 or AIX 5L Version 5.1	**Version 4.2** TCP/IP, which comes standard with AIX Shared Memory Protocol (with AIX server only)	**Version 4.2** The backup-archive client requires: ► X Windows System, X11R6 ► Motif 1.2 or 2.0 ► CDE ► JFS

8.5 Shell script

With AIX 5L, there are a few functional enhancements of the ksh, the Korn shell. The default shell is still /usr/bin/ksh. It is hard linked to /usr/bin/psh, /usr/bin/sh, and /usr/bin/tsh. It is an implementation of the 1988 version of the Korn Shell, enhanced to be POSIX compliant. For a detailed description of the enhancements, refer to the official KornShell Web site at:

http://www.kornshell.com

One important issue is that the shell attribute is changed to the root user. In AIX Version 4.3, it is /bin/ksh, which means it relies on the link between /bin and /usr/bin. This could lead to boot problems because of the unavailability of a shell. The new shell attribute is therefore /usr/bin/ksh. This means that as part of your migration task, you need to verify the compatibility of your shell scripts.

In Example 8-6 on page 195, we show an example where some of the enhancements of the korn shell functionality has a great impact on the application. If you look at the results below, you can see that we get an error message or even a wrong calculated result when we do the same calculation with AIX 5L as we did with AIX Version 4.3.

Example 8-6 Shell changes

```
#In AIX Version 4.3#
# echo $((07-1))
6
# echo $((08-1))
7
# echo $((10-1))
9
# echo $((010-1))
9
#In AIX Version 5L (with APAR IY22337)#
# echo $((07-1))
6
# echo $((08-1))
ksh: 08-1: 0403-009 The specified number is not valid for this command.
# echo $((10-1))
9
# echo $((010-1))
7
```

The reason for this error is that the behavior of `$((...))` was changed with APAR IY22337 to bring the korn shell into conformance with the POSIX standard. That standard states that a numeric expression beginning with a leading '0' will be interpreted as an *octal* number. In AIX Version 4.3, however a number beginning with '0' is considered to be a *decimal* number. This has an major impact on scripts that calculate dates or duration numbers in terms of minutes or seconds. We therefore recommend that if you use scripts in your environment on a regular basis, you need to study the latest ksh enhancements in detail to ensure they are still working as desired after your migration process.

8.6 Compiler

This information on compiler support for AIX 5L Version 5.1 and Version 5.2 is taken from the *AIX 5L for POWER Version 5.1 Release Notes*, GI10-0729 or *AIX 5L Version 5.2 Release Notes*, GI10-0739. If the compiler issue is of any concern to you, verify that this information is still valid. There may be additional versions supported already. The information we give in this section is valid at time of writing this book.

AIX 5L Version 5.1

Support for the AIX 5L Version 5.1 64-bit large datatype and execution on both 32-bit and 64-bit kernels will be available via APARs IY16228 and IY16948.

The following compilers are fully supported with AIX 5L Version 5.1:

- VisualAge® C++ Professional for AIX, Version 6.0.0
- VisualAge C++ Professional for AIX, Version 5.0.2
- C for AIX, Version 6.0.0
- C for AIX, Version 5.0.2
- XL Fortran for AIX, Version 8.1
- XL Fortran for AIX, Version 7.1.1

AIX 5L Version 5.2

Similarly, for AIX 5L Version 5.2, we give a list of supported compilers:

- VisualAge C++ Professional for AIX, Version 6.0.0. You need to install Version 6.0.0 as well as APARs IY34533, IY34534, IY34536, IY34538, and IY34623.
- VisualAge C++ Professional for AIX, Version 5.0.2. You need to install Version 5.0.2 as well as APARs IY34533, IY34534, IY34535, IY34537, and IY34623.
- C for AIX, Version 6.0.0. You are required to install Version 6.0.0 and APARs IY34533, IY34534, IY34536, and IY34623.
- C for AIX, Version 5.0.2. You are required to install Version 5 and then apply APARs IY34533, IY34534, IY34535, and IY34623.
- XL Fortran for AIX, Version 8.1. Similarly, install Version 8.1 and apply APARs IY34533, IY34534, IY33757, IY33758, and IY34623.
- XL Fortran for AIX, Version 7.1.1. Install Version 7.1.1 and additionally APARs IY34533, IY34534, IY33755, IY33756, and IY34623.

A

Hardware and fileset support with AIX 5L

We give a detailed list of supported and unsupported hardware components in this appendix. The first section involves AIX 5L Version 5.1. The next section involves supported and unsupported hardware with AIX 5L Version 5.2.

AIX 5L Version 5.1

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

- RS/6000 7013 Models S70 and S7A
- RS/6000 7015 Models S70 and S7A
- RS/6000 7017 Models S70, S7A, and S80
- RS/6000 7025 Models H80 and F80
- RS/6000 7026 Models H70, H80, and M80
- RS/6000 7043 Models 260 and 270
- RS/6000 7044 Models 170 and 270
- IBM @server pSeries 680 Model S85
- IBM @server pSeries 640 Model B80
- IBM @server pSeries 660 Model 6H1
- IBM @server pSeries 660 Model 6M1
- IBM @server pSeries 620 Model 6F1
- IBM @server pSeries 690
- IBM @server pSeries 610 Model 6C1
- IBM @server pSeries 610 Model 6E1
- IBM @server pSeries 670
- IBM @server pSeries 630 Model 6C4
- IBM @server pSeries 630 Model 6E4
- IBM @server pSeries 650
- IBM @server pSeries 655

Unsupported devices and machines

- AIX Ultimedia Services Audio and Video devices
- PCMCIA device support
- 2751 ESCON® CNTRL UNIT, PCI/LONG/32BIT/5V
- 2947 ARTIC960HX ADPTR + 4PRT MP AIB, PCI/LONG/32BIT/3.3, or 5V
- 6310 ARTIC960RXD, DIGITAL TRUNK, PCI/LONG/32BIT/3.3, or 5V

Unsupported functions and filesets

- 7318 Model P10/S20 Serial Communications Network Server
- AIX Xstation Manager®
- AIX Version 3.2 Network Installation Tools
- Remote Customer Support and Services
- SOMobjects® Base Toolkit
- Information Presentation Facility Runtime
- X11.vsm.helps

- X11.vsm.icons
- X11.vsm.rte
- libipfx.a
- The 7318 Serial Communications Network Server, which provides serial and parallel connectivity to Ethernet Networks
- Network Terminal Accelerator
- The 9333 Serial Link DASD Subsystem
- devices.pci.b7105090

AIX 5L Version 5.2

For supported systems, refer to the list we included in “AIX 5L Version 5.1” on page 198. For unsupported devices and machines, refer to the following section.

Unsupported devices and machines

- RS/6000 or OEM hardware based on the MCA bus
- Scalable Parallel (SP) nodes based on the MCA bus
- RS/6000, Power Personal Systems, or OEM hardware based on the PReP architecture
- POWER1, POWER2, POWER Singe Chip (RSC), POWER2 Single Chip (P2RSC), or 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 (GTX 150P) PCI/SHORT/32BIT/5V, GRAPHICS ADAPTER
  - 2657 S15 GRAPHICS ADAPTER, PCI/SHORT/32BIT/5V, WEITEK P9100
  - 2837 MVP MULTI-MONITOR ADAPTER, PCI/LONG/32BIT/3.3 or 5V
  - 2854 3D (GTX500P), PCI/LONG/32BIT/3.3 or 5V, GRAPHICS ADAPTER
  - 2855 3DX (GTX550P), PCI/LONG/32BIT/3.3 or 5V, GRAPHICS ADAPTER
  - 2856 PCI/SHORT/32-BIT/3.3 or 5V, 7250 ATTACH ADAPTER
  - 8242 10/100BASET ETHERNET PCI/SHORT/32BIT/5V

- 2751 ESCON CNTRL UNIT, PCI/LONG/32BIT/5V
- 2947 ARTIC960HX ADPTR + 4PRT MP AIB, PCI/LONG/32BIT/3.3 or 5V
- 6310 ARTIC960RXD, DIGITAL TRUNK, PCI/LONG/32BIT/3.3 or 5V

- ISA adapters:
  - 2647 VIDEO CAPTURE ENHANCEMENT, ISA/SHORT
  - 2701 4 PORT SDLC, ISA/LONG, EIA 232/V.35/X.21
  - 2931 8-PORT, ISA/LONG, EIA232 ADAPTER/FAN-OUT BOX
  - 2932 8-PORT, ISA/LONG, EIA232/422 ADAPTER/FAN-OUT BOX
  - 2933 128-PORT, ISA/LONG, EIA232 ASYNCH CONTROLLER
  - 2961 1 PORT X.25, SDLC, PPP, ISA/LONG, ADAPTER (C1X)
  - 2971 TOKEN RING ADAPTER, ISA
  - 2981 ETHERNET ADAPTER, ISA, RJ45/BNC
  - 8240 A/M 3COM ETHERNET ISA/SHORT TP ONLY
  - 8241 A/M 3COM ETHERNET ISA/SHORT BNC/AUI
- Non-CHRP graphics adapters:
  - Gt3/Gt3i
  - Gt4/Gt4e/Gt4i/Gt4x/Gt4xi
  - GXT110P
  - GXT150L™/GXT150M™/GXT150P
  - GXT155L
  - GXT500
  - GXT500D
  - GXT500P
  - GXT550P (feature code (FC) 2855 only)
  - GXT800M
  - GXT1000™
  - MVP MULTIPCI Adapter
  - S15
  - VIDEO OUTPUT OPTION (#3200) (FC 7254)
  - 7250 ATTACH Adapter (FC 2856)

Unsupported functions and filesets

- 7318 Model p10/S20 Serial Communications Network Server
- AIX Xstation Manager
- AIX Version 3.2 Network Installation Tools
- Remote Customer Support and Services
- SOMobjects Base Toolkit
- Information Presentation Facility Runtime
- X11.vsm.helps
- X11.vsm.icons
- X11.vsm.rte
- GL 3.2
- Power management
- IBM-850 locales
- libipfx.a
- devices.pci.b7105090
- The 7318 Serial Communications Network Server
- Network Terminal Accelerator
- The 9333 Serial Link DASD Subsystem

B

Sample output of the snap command

In the first example, Example B-1, we give an output of the file system information collected from running the **snap** command. This information is collected in the filesys.snap file.

Example: B-1 Information collected in filesys.snap

```
# view /tmp/ibmsupt/filesys/filesys.snap
.....
.....    creation date
.....

Wed Jan 29 16:35:30 CST 2003

.....
.....    df -k
.....

Filesystem    1024-blocks      Free %Used    Iused %Iused Mounted on
/dev/hd4            16384      8780   47%     1098    14% /
/dev/hd2           507904     44336   92%    15365    13% /usr
/dev/hd9var         16384     14440   12%      155     4% /var
/dev/hd3           114688    106428    8%      112     1% /tmp
/dev/hd1            16384     15820    4%       20     1% /home
```

```
.....
.....    mount
.....

  node       mounted        mounted over    vfs       date        options
-------- ---------------  ---------------  ------ ------------ ---------------
         /dev/hd4         /                jfs    Jan 14 18:22 rw,log=/dev/hd8
         /dev/hd2         /usr             jfs    Jan 14 18:22 rw,log=/dev/hd8
         /dev/hd9var      /var             jfs    Jan 14 18:22 rw,log=/dev/hd8
         /dev/hd3         /tmp             jfs    Jan 14 18:22 rw,log=/dev/hd8
         /dev/hd1         /home            jfs    Jan 14 18:24 rw,log=/dev/hd8

.....
.....    lsfs -l
.....

Name            Nodename   Mount Pt               VFS   Size    Options    Auto
Accounting
/dev/hd4        --         /                      jfs   32768   --         yes
no
/dev/hd1        --         /home                  jfs   32768   --         yes
no
/dev/hd2        --         /usr                   jfs   1015808 --         yes
no
/dev/hd9var     --         /var                   jfs   32768   --         yes
no
/dev/hd3        --         /tmp                   jfs   229376  --         yes
no

.....
.....    lsvg
.....

rootvg

.....
.....    lsvg -p rootvg
.....

rootvg:
PV_NAME           PV STATE    TOTAL PPs   FREE PPs    FREE DISTRIBUTION
hdisk0            active      542         371         108..00..46..108..109

.....
.....    lsvg | xargs lsvg -l
.....

rootvg:
LV NAME             TYPE       LPs   PPs   PVs  LV STATE      MOUNT POINT
```

```
hd5                 boot       1     1     1    closed/syncd  N/A
hd6                 paging     128   128   1    open/syncd    N/A
hd8                 jfslog     1     1     1    open/syncd    N/A
hd4                 jfs        1     1     1    open/syncd    /
hd2                 jfs        31    31    1    open/syncd    /usr
hd9var              jfs        1     1     1    open/syncd    /var
hd3                 jfs        7     7     1    open/syncd    /tmp
hd1                 jfs        1     1     1    open/syncd    /home

.....
.....    lspv -l hdisk0
.....

hdisk0:
LV NAME               LPs   PPs   DISTRIBUTION          MOUNT POINT
hd5                   1     1     01..00..00..00..00    N/A
hd6                   128   128   00..108..20..00..00   N/A
hd8                   1     1     00..00..01..00..00    N/A
hd4                   1     1     00..00..01..00..00    /
hd2                   31    31    00..00..31..00..00    /usr
hd9var                1     1     00..00..01..00..00    /var
hd3                   7     7     00..00..07..00..00    /tmp
hd1                   1     1     00..00..01..00..00    /home

.....
.....    lspv -l hdisk1
.....

0516-320 : Physical volume 0001615fcbc1a83f0000000000000000 is not assigned to
        a volume group.

.....
.....    lspv -l hdisk2
.....

0516-320 : Physical volume 0001615fcbc1a86b0000000000000000 is not assigned to
        a volume group.

.....
.....    lspv -l hdisk3
.....

0516-320 : Physical volume 0001615fcbea5d160000000000000000 is not assigned to
        a volume group.

.....
.....    lspv -l hdisk4
.....
```

```
0516-320 : Physical volume 0001615fcbea5e580000000000000000 is not assigned to
        a volume group.

.....
.....    lspv -l hdisk5
.....

0516-320 : Physical volume 0001615fcbea5f960000000000000000 is not assigned to
        a volume group.

.....
.....    lslv -l hd5
.....

hd5:N/A
PV                COPIES        IN BAND       DISTRIBUTION
hdisk0            001:000:000   100%          001:000:000:000:000

.....
.....    lslv -l hd6
.....

hd6:N/A
PV                COPIES        IN BAND       DISTRIBUTION
hdisk0            128:000:000   84%           000:108:020:000:000

.....
.....    lslv -l hd8
.....

hd8:N/A
PV                COPIES        IN BAND       DISTRIBUTION
hdisk0            001:000:000   100%          000:000:001:000:000

.....
.....    lslv -l hd4
.....

hd4:/
PV                COPIES        IN BAND       DISTRIBUTION
hdisk0            001:000:000   100%          000:000:001:000:000

.....
.....    lslv -l hd2
.....

hd2:/usr
PV                COPIES        IN BAND       DISTRIBUTION
hdisk0            031:000:000   100%          000:000:031:000:000
```

```
.....
.....    lslv -l hd9var
.....

hd9var:/var
PV               COPIES        IN BAND       DISTRIBUTION
hdisk0           001:000:000   100%          000:000:001:000:000

.....
.....    lslv -l hd3
.....

hd3:/tmp
PV               COPIES        IN BAND       DISTRIBUTION
hdisk0           007:000:000   100%          000:000:007:000:000

.....
.....    lslv -l hd1
.....

hd1:/home
PV               COPIES        IN BAND       DISTRIBUTION
hdisk0           001:000:000   100%          000:000:001:000:000

.....
.....    lsattr -El hdisk0
.....

pvid        0001615f43bd81800000000000000000 Physical volume identifier False
queue_depth 3                                Queue DEPTH                False
size_in_mb  9100                             Size in Megabytes          False

.....
.....    lsattr -El hdisk1
.....

pvid        0001615fcbc1a83f0000000000000000 Physical volume identifier False
queue_depth 3                                Queue DEPTH                False
size_in_mb  9100                             Size in Megabytes          False

.....
.....    lsattr -El hdisk2
.....

pvid        0001615fcbc1a86b0000000000000000 Physical volume identifier False
queue_depth 3                                Queue DEPTH                False
size_in_mb  9100                             Size in Megabytes          False
```

```
.....
.....    lsattr -El hdisk3
.....

pvid        0001615fcbea5d160000000000000000 Physical volume identifier False
queue_depth 3                                Queue DEPTH                False
size_in_mb  9100                             Size in Megabytes          False

.....
.....    lsattr -El hdisk4
.....

pvid        0001615fcbea5e580000000000000000 Physical volume identifier False
queue_depth 3                                Queue DEPTH                False
size_in_mb  9100                             Size in Megabytes          False

.....
.....    lsattr -El hdisk5
.....

pvid        0001615fcbea5f960000000000000000 Physical volume identifier False
queue_depth 3                                Queue DEPTH                False
size_in_mb  9100                             Size in Megabytes          False
```

Similarly, in Example B-2 we show the output collected after running the **snap** command, which is collected in the tcpip.snap file.

Example: B-2 Information collected in tcpip.snap

```
.....
.....    Note: If you want the /var/adm/ras/trcfile, snap -g will get it
.....


.....
.....    lssrc -a
.....

Subsystem         Group            PID     Status
 syslogd          ras              6196    active
 sendmail         mail             7256    active
 portmap          portmap          7510    active
 inetd            tcpip            8304    active
 snmpd            tcpip            7774    active
 dpid2            tcpip            8010    active
 biod             nfs              9808    active
 rpc.statd        nfs              9554    active
 rpc.lockd        nfs              8794    active
```

```
qdaemon          spooler          10582   active
writesrv         spooler                  inoperative
lpd              spooler                  inoperative
clvmd                                     inoperative
gated            tcpip                    inoperative
named            tcpip                    inoperative
routed           tcpip                    inoperative
rwhod            tcpip                    inoperative
iptrace          tcpip                    inoperative
xntpd            tcpip                    inoperative
timed            tcpip                    inoperative
dhcpcd           tcpip                    inoperative
dhcpsd           tcpip                    inoperative
dhcprd           tcpip                    inoperative
ndpd-host        tcpip                    inoperative
ndpd-router      tcpip                    inoperative
tftpd            tcpip                    inoperative
nfsd             nfs                      inoperative
rpc.mountd       nfs                      inoperative
automountd       autofs                   inoperative
llbd             iforncs                  inoperative
glbd             iforncs                  inoperative
i4lmd            iforls                   inoperative
i4glbcd          iforncs                  inoperative
i4gdb            iforls                   inoperative
i4llmd           iforls                   inoperative
wsmrefserver                              inoperative
dtsrc                                     inoperative

.....
.....    netstat -m
.....


Kernel malloc statistics:

******* CPU 0 *******
By size       inuse     calls failed   delayed    free   hiwat   freed
32              100      6473      0         0      28    1440       0
64               38       200      0         0      26     720       0
128              25      2947      0         0       7     360       0
256             371     96237      3         0     525     864       0
512              27      5643      0         0       5      90       0
1024             12      4348      0         0       8     225       0
2048            384      3476     20         0     224     225     870
4096              2       271      0         0     146     270       0
8192              3         9      0         0       0      22       0
16384             0         0      0         0      40      54       0
65536             1         1      0         0       0    2047       0
```

```
******* CPU 1 *******
By size       inuse     calls failed   delayed    free   hiwat   freed
32               42      5265      0         0      86    1440       0
64                7        72      0         0      57     720       0
128              17      1933      0         0      47     360       0
256              15     80468      0         0     513     864       0
512              26      4890      0         0       6      90       0
1024             14      3383      0         0       6     225       0
2048              0      1413      4         0     224     225     370
4096              1       202      0         0     144     270       0
16384             1         1      0         0      40      54       0


******* CPU 2 *******
By size       inuse     calls failed   delayed    free   hiwat   freed
32               29      6632      0         0      99    1440       0
64               13        54      0         0      51     720       0
128               8      2573      0         0      24     360       0
256              18     99660      0         0     510     864       0
512              15      5721      0         0       9      90       0
1024             21      5052      0         0       3     225       0
2048              0       517      0         0     224     225     145
4096              0        66      0         0     143     270       0
8192              1         2      0         0       0      22       0
16384             0         0      0         0      40      54       0


******* CPU 3 *******
By size       inuse     calls failed   delayed    free   hiwat   freed
32               43      5880      0         0      85    1440       0
64               29        57      0         0      35     720       0
128              25      2771      0         0      39     360       0
256              27     96428      0         0     517     864       0
512              16      5055      0         0      16      90       0
1024             20      5066      0         0       8     225       0
2048              0       513      5         0     224     225     144
4096              0         5      0         0     144     270       0
8192              0         1      0         0       0      22       0
16384             0         0      0         0      40      54       0

By type       inuse   calls failed delayed  memuse  memmax  mapb

Streams mblk statistic failures:
0 high priority mblk failures
0 medium priority mblk failures
0 low priority mblk failures
```

```
.....
.....    netstat -in
.....

Name  Mtu   Network     Address              Ipkts Ierrs    Opkts Oerrs  Coll
en1   1500  link#2      0.6.29.b9.1e.ae    7319323     0    22246     0     0
en1   1500  9.3.4       9.3.4.34           7319323     0    22246     0     0
lo0   16896 link#1                          137173     0   137176     0     0
lo0   16896 127         127.0.0.1          137173     0   137176     0     0
lo0   16896 ::1                            137173     0   137176     0     0

.....
.....    netstat -v
.....

-------------------------------------------------------------
ETHERNET STATISTICS (ent0) :
Device Type: IBM 10/100/1000 Base-T Ethernet PCI Adapter (14100401)
Hardware Address: 00:06:29:6b:0d:3d
Elapsed Time: 0 days 0 hours 0 minutes 2 seconds

Transmit Statistics:                          Receive Statistics:
--------------------                          -------------------
Packets: 0                                    Packets: 0
Bytes: 0                                      Bytes: 0
Interrupts: 0                                 Interrupts: 0
Transmit Errors: 0                            Receive Errors: 0
Packets Dropped: 0                            Packets Dropped: 0
                                              Bad Packets: 0
Max Packets on S/W Transmit Queue: 0
S/W Transmit Queue Overflow: 0
Current S/W+H/W Transmit Queue Length: 0

Broadcast Packets: 0                          Broadcast Packets: 0
Multicast Packets: 0                          Multicast Packets: 0
No Carrier Sense: 0                           CRC Errors: 0
DMA Underrun: 0                               DMA Overrun: 0
Lost CTS Errors: 0                            Alignment Errors: 0
Max Collision Errors: 0                       No Resource Errors: 0
Late Collision Errors: 0                      Receive Collision Errors: 0
Deferred: 0                                   Packet Too Short Errors: 0
SQE Test: 0                                   Packet Too Long Errors: 0
Timeout Errors: 0                             Packets Discarded by Adapter: 0
Single Collision Count: 0                     Receiver Start Count: 0
Multiple Collision Count: 0
Current HW Transmit Queue Length: 0

General Statistics:
-------------------
```

```
No mbuf Errors: 0
Adapter Reset Count: 0
Adapter Data Rate: 0
Driver Flags: Up Broadcast Simplex
        Limbo 64BitSupport ChecksumTCP
        ChecksumOffload PrivateSegment DataRateSet

Adapter Specific Statistics:
----------------------------
Additional Driver Flags: Autonegotiate
Entries to transmit timeout routine: 0
Firmware Level: 12.4.15
Transmit and Receive Flow Control Status: Disabled
Link Status: Down
Media Speed Selected: Autonegotiation
Media Speed Running: Unknown
Packets with Transmit collisions:
 1 collisions: 0          6 collisions: 0          11 collisions: 0
 2 collisions: 0          7 collisions: 0          12 collisions: 0
 3 collisions: 0          8 collisions: 0          13 collisions: 0
 4 collisions: 0          9 collisions: 0          14 collisions: 0
 5 collisions: 0         10 collisions: 0          15 collisions: 0
-------------------------------------------------------------
TOKENRING STATISTICS (tok0) :
Device Type: IBM PCI Tokenring Adapter (14103e00)
Hardware Address: 00:60:94:8a:d3:58
Elapsed Time: 0 days 0 hours 0 minutes 26 seconds

Transmit Statistics:                          Receive Statistics:
--------------------                          -------------------
Packets: 0                                    Packets: 0
Bytes: 0                                      Bytes: 0
Interrupts: 0                                 Interrupts: 0
Transmit Errors: 0                            Receive Errors: 0
Packets Dropped: 0                            Packets Dropped: 0
                                              Bad Packets: 0
Max Packets on S/W Transmit Queue: 0
S/W Transmit Queue Overflow: 0
Current S/W+H/W Transmit Queue Length: 0

Broadcast Packets: 0                          Broadcast Packets: 0
Multicast Packets: 0                          Multicast Packets: 0
Timeout Errors: 0                             Receive Congestion Errors: 0
Current SW Transmit Queue Length: 0
Current HW Transmit Queue Length: 0

General Statistics:
-------------------
No mbuf Errors: 0                             Lobe Wire Faults: 0
```

```
Abort Errors: 0                                 AC Errors: 0
Burst Errors: 0                                 Frame Copy Errors: 0
Frequency Errors: 0                             Hard Errors: 0
Internal Errors: 0                              Line Errors: 0
Lost Frame Errors: 0                            Only Station: 0
Token Errors: 0                                 Remove Received: 0
Ring Recovered: 0                               Signal Loss Errors: 0
Soft Errors: 0                                  Transmit Beacon Errors: 0
Driver Flags: Up Broadcast Limbo
        64BitSupport

IBM PCI Tokenring Adapter (14103e00) Specific Statistics:
---------------------------------------------------------
Media Speed Running: 4 Mbps Half Duplex
Media Speed Selected: Auto negotiation / Full Duplex
Receive Overruns : 0
Transmit Underruns : 0
ARI/FCI Errors : 0
Microcode level on the adapter : 00IHSS2B4
Num pkts in priority sw tx queue : 0
Num pkts in priority hw tx queue : 0
Open Firware level : 001PXHL00
---------------------------------------------------------
ETHERNET STATISTICS (ent1) :
Device Type: IBM 10/100 Mbps Ethernet PCI Adapter (23100020)
Hardware Address: 00:06:29:b9:1e:ae
Elapsed Time: 12 days 2 hours 37 minutes 6 seconds

Transmit Statistics:                            Receive Statistics:
--------------------                            -------------------
Packets: 22246                                  Packets: 7823295
Bytes: 1929897                                  Bytes: 2589386009
Interrupts: 8                                   Interrupts: 7823289
Transmit Errors: 0                              Receive Errors: 0
Packets Dropped: 1                              Packets Dropped: 1
                                                Bad Packets: 0
Max Packets on S/W Transmit Queue: 7
S/W Transmit Queue Overflow: 0
Current S/W+H/W Transmit Queue Length: 0

Broadcast Packets: 3                            Broadcast Packets: 2074985
Multicast Packets: 2                            Multicast Packets: 5732348
No Carrier Sense: 0                             CRC Errors: 0
DMA Underrun: 0                                 DMA Overrun: 0
Lost CTS Errors: 0                              Alignment Errors: 0
Max Collision Errors: 0                         No Resource Errors: 0
Late Collision Errors: 0                        Receive Collision Errors: 0
Deferred: 2220                                  Packet Too Short Errors: 0
SQE Test: 0                                     Packet Too Long Errors: 0
```

```
Timeout Errors: 0                               Packets Discarded by Adapter: 0
Single Collision Count: 435                     Receiver Start Count: 0
Multiple Collision Count: 249
Current HW Transmit Queue Length: 0

General Statistics:
-------------------
No mbuf Errors: 0
Adapter Reset Count: 1
Adapter Data Rate: 10
Driver Flags: Up Broadcast Running
        Simplex AlternateAddress 64BitSupport
        PrivateSegment DataRateSet

IBM 10/100 Mbps Ethernet PCI Adapter (23100020) Specific Statistics:
------------------------------------------------
Chip Version: 26
RJ45 Port Link Status : up
Media Speed Selected: Auto negotiation
Media Speed Running: 10 Mbps Half Duplex
Receive Pool Buffer Size: 384
Free Receive Pool Buffers: 128
No Receive Pool Buffer Errors: 0
Inter Packet Gap: 96
Adapter Restarts due to IOCTL commands: 1
Packets with Transmit collisions:
 1 collisions: 435         6 collisions: 0          11 collisions: 0
 2 collisions: 227         7 collisions: 0          12 collisions: 0
 3 collisions: 20          8 collisions: 0          13 collisions: 0
 4 collisions: 1           9 collisions: 0          14 collisions: 0
 5 collisions: 1          10 collisions: 0          15 collisions: 0
Excessive Deferrals: 0
------------------------------------------------------------
ETHERNET STATISTICS (ent2) :
Device Type: IBM 10/100/1000 Base-T Ethernet PCI Adapter (14100401)
Hardware Address: 00:06:29:6b:0d:32
Elapsed Time: 0 days 0 hours 0 minutes 2 seconds

Transmit Statistics:                            Receive Statistics:
--------------------                            -------------------
Packets: 0                                      Packets: 0
Bytes: 0                                        Bytes: 0
Interrupts: 0                                   Interrupts: 0
Transmit Errors: 0                              Receive Errors: 0
Packets Dropped: 0                              Packets Dropped: 0
                                                Bad Packets: 0

Max Packets on S/W Transmit Queue: 0
S/W Transmit Queue Overflow: 0
Current S/W+H/W Transmit Queue Length: 0
```

```
Broadcast Packets: 0                        Broadcast Packets: 0
Multicast Packets: 0                        Multicast Packets: 0
No Carrier Sense: 0                         CRC Errors: 0
DMA Underrun: 0                             DMA Overrun: 0
Lost CTS Errors: 0                          Alignment Errors: 0
Max Collision Errors: 0                     No Resource Errors: 0
Late Collision Errors: 0                    Receive Collision Errors: 0
Deferred: 0                                 Packet Too Short Errors: 0
SQE Test: 0                                 Packet Too Long Errors: 0
Timeout Errors: 0                           Packets Discarded by Adapter: 0
Single Collision Count: 0                   Receiver Start Count: 0
Multiple Collision Count: 0
Current HW Transmit Queue Length: 0

General Statistics:
-------------------
No mbuf Errors: 0
Adapter Reset Count: 0
Adapter Data Rate: 0
Driver Flags: Up Broadcast Simplex
        Limbo 64BitSupport ChecksumTCP
        ChecksumOffload PrivateSegment DataRateSet

Adapter Specific Statistics:
----------------------------
Additional Driver Flags: Autonegotiate
Entries to transmit timeout routine: 0
Firmware Level: 12.4.15
Transmit and Receive Flow Control Status: Disabled
Link Status: Down
Media Speed Selected: Autonegotiation
Media Speed Running: Unknown
Packets with Transmit collisions:
 1 collisions: 0           6 collisions: 0          11 collisions: 0
 2 collisions: 0           7 collisions: 0          12 collisions: 0
 3 collisions: 0           8 collisions: 0          13 collisions: 0
 4 collisions: 0           9 collisions: 0          14 collisions: 0
 5 collisions: 0          10 collisions: 0          15 collisions: 0

.....
.....    netstat -s
.....

ip:
        6414631 total packets received
        0 bad header checksums
        0 with size smaller than minimum
        0 with data size < data length
```

```
        0 with header length < data size
        0 with data length < header length
        0 with bad options
        0 with incorrect version number
        0 fragments received
        0 fragments dropped (dup or out of space)
        0 fragments dropped after timeout
        0 packets reassembled ok
        658989 packets for this host
        121 packets for unknown/unsupported protocol
        0 packets forwarded
        23179 packets not forwardable
        0 redirects sent
        152648 packets sent from this host
        0 packets sent with fabricated ip header
        0 output packets dropped due to no bufs, etc.
        0 output packets discarded due to no route
        0 output datagrams fragmented
        0 fragments created
        0 datagrams that can't be fragmented
        5732342 IP Multicast packets dropped due to no receiver
        0 successful path MTU discovery cycles
        0 path MTU rediscovery cycles attempted
        0 path MTU discovery no-response estimates
        0 path MTU discovery response timeouts
        0 path MTU discovery decreases detected
        0 path MTU discovery packets sent
        0 path MTU discovery memory allocation failures
        0 ipintrq overflows
        0 with illegal source
        0 packets processed by threads
        0 packets dropped by threads
icmp:
        123 calls to icmp_error
        0 errors not generated because old message was icmp
        Output histogram:
                echo reply: 6569
                destination unreachable: 123
        0 messages with bad code fields
        0 messages < minimum length
        0 bad checksums
        0 messages with bad length
        Input histogram:
                echo reply: 2
                destination unreachable: 115
                echo: 6571
        6569 message responses generated
igmp:
        6 messages received
```

```
        0 messages received with too few bytes
        0 messages received with bad checksum
        0 membership queries received
        0 membership queries received with invalid field(s)
        6 membership reports received
        0 membership reports received with invalid field(s)
        6 membership reports received for groups to which we belong
        2 membership reports sent
tcp:
        132096 packets sent
                77508 data packets (10141254 bytes)
                0 data packets (0 bytes) retransmitted
                37000 ack-only packets (13384 delayed)
                0 URG only packets
                0 window probe packets
                0 window update packets
                17588 control packets
        132709 packets received
                84930 acks (for 10158809 bytes)
                6028 duplicate acks
                0 acks for unsent data
                82785 packets (10069253 bytes) received in-sequence
                167 completely duplicate packets (167 bytes)
                0 old duplicate packets
                0 packets with some dup. data (0 bytes duped)
                5856 out-of-order packets (0 bytes)
                0 packets (0 bytes) of data after window
                0 window probes
                4 window update packets
                1 packet received after close
                0 packets with bad hardware assisted checksum
                0 discarded for bad checksums
                0 discarded for bad header offset fields
                0 discarded because packet too short
        5881 connection requests
        5874 connection accepts
        11742 connections established (including accepts)
        14659 connections closed (including 1 drop)
        13 embryonic connections dropped
        90796 segments updated rtt (of 90810 attempts)
        0 resends due to path MTU discovery
        0 path MTU discovery terminations due to retransmits
        0 retransmit timeouts
                0 connections dropped by rexmit timeout
        0 persist timeouts
        1587 keepalive timeouts
                167 keepalive probes sent
        0 connections dropped by keepalive
        0 times SACK blocks array is extended
```

```
        0 times SACK holes array is extended
        0 packets dropped due to memory allocation failure
        0 connections in timewait reused
        0 delayed ACKs for SYN
        0 delayed ACKs for FIN
        0 send_and_disconnects
udp:
        519707 datagrams received
        0 incomplete headers
        0 bad data length fields
        0 bad checksums
        123 dropped due to no socket
        506005 broadcast/multicast datagrams dropped due to no socket
        0 dropped due to full socket buffers
        13579 delivered
        13658 datagrams output

ipv6:
        0 total packets received
        0 with size smaller than minimum
        0 with data size < data length
        0 with incorrect version number
        0 with illegal source
        0 input packets without enough memory
        0 fragments received
        0 fragments dropped (dup or out of space)
        0 fragments dropped after timeout
        0 packets reassembled ok
        0 packets for this host
        0 packets for unknown/unsupported protocol
        0 packets forwarded
        0 packets not forwardable
        0 too big packets not forwarded
        0 packets sent from this host
        0 packets sent with fabricated ipv6 header
        0 output packets dropped due to no bufs
        0 output packets without enough memory
        0 output packets discarded due to no route
        0 output datagrams fragmented
        0 fragments created
icmpv6:
        0 calls to icmp6_error
        0 errors not generated because old message was icmpv6
        Output histogram:
        unreachables: 0
                packet too bigs: 0
                time exceededs: 0
                parameter problems: 0
                redirects: 0
```

```
                echo requests: 0
                echo replies: 0
                group queries: 0
                group reports: 0
                group terminations: 0
                router solicitations: 0
                router advertisements: 0
                neighbor solicitations: 0
                neighbor advertisements: 0
        0 messages with bad code fields
        0 messages < minimum length
        0 bad checksums
        0 messages with bad length
        Input histogram:
                unreachables: 0
                packet too bigs: 0
                time exceededs: 0
                parameter problems: 0
                echo requests: 0
                echo replies: 0
                group queries: 0
                        bad group queries: 0
                group reports: 0
                        bad group reports: 0
                        our groups' reports: 0
                group terminations: 0
                bad group terminations: 0
                router solicitations: 0
                bad router solicitations: 0
                router advertisements: 0
                bad router advertisements: 0
                neighbor solicitations: 0
                bad neighbor solicitations: 0
                neighbor advertisements: 0
                bad neighbor advertisements: 0
                redirects: 0
                bad redirects: 0
        0 message responses generated

.....
.....    netstat -an
.....

Active Internet connections (including servers)
Proto Recv-Q Send-Q  Local Address          Foreign Address        (state)
tcp4       0      0  9.3.4.34.38644         9.3.4.34.38645         ESTABLISHED
tcp4       0      0  9.3.4.34.38645         9.3.4.34.38644         ESTABLISHED
tcp4       0      0  *.38644                *.*                    LISTEN
tcp4       0      0  9.3.4.34.32786         9.3.4.34.777           ESTABLISHED
```

```
tcp4       0      0  9.3.4.34.777           9.3.4.34.32786         ESTABLISHED
tcp4       0      0  9.3.4.34.23            9.3.5.15.1598          ESTABLISHED
tcp4       0      0  9.3.4.34.23            9.3.5.15.1549          ESTABLISHED
tcp4       0      0  9.3.4.34.32803         9.3.4.34.32804         ESTABLISHED
tcp4       0      0  9.3.4.34.32804         9.3.4.34.32803         ESTABLISHED
tcp4       0      0  *.32803                *.*                    LISTEN
tcp4       0      0  9.3.4.34.32786         9.3.4.34.653           ESTABLISHED
tcp4       0      0  9.3.4.34.653           9.3.4.34.32786         ESTABLISHED
tcp4       0      0  9.3.4.34.32795         9.3.4.34.32796         ESTABLISHED
tcp4       0      0  9.3.4.34.32796         9.3.4.34.32795         ESTABLISHED
tcp4       0      0  *.32795                *.*                    LISTEN
tcp4       0      0  9.3.4.34.32786         9.3.4.34.924           ESTABLISHED
tcp4       0      0  9.3.4.34.924           9.3.4.34.32786         ESTABLISHED
tcp4       0      0  9.3.4.34.32789         9.3.4.34.32790         ESTABLISHED
tcp4       0      0  9.3.4.34.32790         9.3.4.34.32789         ESTABLISHED
tcp4       0      0  *.32789                *.*                    LISTEN
tcp4       0      0  9.3.4.34.32786         9.3.4.34.651           ESTABLISHED
tcp4       0      0  9.3.4.34.651           9.3.4.34.32786         ESTABLISHED
tcp4       0      0  9.3.4.34.32787         9.3.4.34.32788         ESTABLISHED
tcp4       0      0  9.3.4.34.32788         9.3.4.34.32787         ESTABLISHED
tcp4       0      0  *.32787                *.*                    LISTEN
tcp4       0      0  9.3.4.34.32786         9.3.4.34.732           ESTABLISHED
tcp4       0      0  9.3.4.34.732           9.3.4.34.32786         ESTABLISHED
tcp4       0      0  9.3.4.34.32769         9.3.4.34.840           ESTABLISHED
tcp4       0      0  9.3.4.34.840           9.3.4.34.32769         ESTABLISHED
tcp4       0      0  9.3.4.34.32769         9.3.4.34.839           ESTABLISHED
tcp4       0      0  9.3.4.34.839           9.3.4.34.32769         ESTABLISHED
tcp4       0      0  9.3.4.34.32769         9.3.4.34.838           ESTABLISHED
tcp4       0      0  9.3.4.34.838           9.3.4.34.32769         ESTABLISHED
tcp4       0      0  9.3.4.34.32769         9.3.4.34.837           ESTABLISHED
tcp4       0      0  9.3.4.34.837           9.3.4.34.32769         ESTABLISHED
tcp4       0      0  *.32786                *.*                    LISTEN
tcp4       0      0  *.6000                 *.*                    LISTEN
tcp4       0      0  *.32785                *.*                    LISTEN
tcp4       0      0  127.0.0.1.199          127.0.0.1.32784        ESTABLISHED
tcp4       0      0  127.0.0.1.32784        127.0.0.1.199          ESTABLISHED
tcp4       0      0  *.199                  *.*                    LISTEN
tcp4       0      0  *.25                   *.*                    LISTEN
tcp4       0      0  127.0.0.1.49213        *.*                    LISTEN
tcp4       0      0  *.32772                *.*                    LISTEN
tcp4       0      0  *.827                  *.*                    LISTEN
tcp4       0      0  *.826                  *.*                    LISTEN
tcp4       0      0  *.6112                 *.*                    LISTEN
tcp4       0      0  *.32769                *.*                    LISTEN
tcp4       0      0  *.37                   *.*                    LISTEN
tcp4       0      0  *.13                   *.*                    LISTEN
tcp4       0      0  *.19                   *.*                    LISTEN
tcp4       0      0  *.9                    *.*                    LISTEN
tcp4       0      0  *.7                    *.*                    LISTEN
```

```
tcp        0      0  *.512                  *.*                    LISTEN
tcp4       0      0  *.543                  *.*                    LISTEN
tcp        0      0  *.513                  *.*                    LISTEN
tcp4       0      0  *.544                  *.*                    LISTEN
tcp        0      0  *.514                  *.*                    LISTEN
tcp        0      0  *.23                   *.*                    LISTEN
tcp        0      0  *.21                   *.*                    LISTEN
tcp4       0      0  *.111                  *.*                    LISTEN
tcp4       0      0  *.32768                *.*                    LISTEN
udp4       0      0  *.7                    *.*
udp4       0      0  *.9                    *.*
udp4       0      0  *.13                   *.*
udp4       0      0  *.19                   *.*
udp4       0      0  *.37                   *.*
udp4       0      0  *.111                  *.*
udp4       0      0  *.161                  *.*
udp4       0      0  *.177                  *.*
udp4       0      0  *.514                  *.*
udp4       0      0  *.518                  *.*
udp4       0      0  *.826                  *.*
udp4       0      0  *.827                  *.*
udp4       0      0  *.32784                *.*
udp4       0      0  *.32785                *.*
udp4       0      0  *.32786                *.*
udp4       0      0  *.32787                *.*
udp4       0      0  *.32788                *.*
udp4       0      0  *.32789                *.*
udp4       0      0  *.32790                *.*
udp4       0      0  *.32800                *.*

Active UNIX domain sockets
SADR/PCB  Type   Recv-Q Send-Q Inode    Conn     Refs     Nextref  Addr
7017ce00 dgram       0      0 13f47928        0        0        0
/dev/.SRC-unix
/SRCxagPEa
7017ef00
7017cc00 dgram       0      0 14175da8        0        0        0
/dev/.SRC-unix
/SRC2ZgPEb
7017eec0
701fac00 dgram       0      0 14d74500        0 701fc7c0        0 /dev/log
701fc800
700fc600 dgram       0      0        0 701fc800        0        0
700b4f00
700fc800 dgram       0      0 139d3f50        0        0        0 /dev/SRC
700b4f40
7017ca00 dgram       0      0 144bb5d0        0        0        0
/dev/.SRC-unix
/SRCBNgPEc
```

```
7017ee80
700bd600 dgram       0      0        0 701fc800        0 700b4f00
700b4dc0
7017c800 dgram       0      0 146e9a50        0        0        0
/dev/.SRC-unix
/SRCGFgPEd
7017ee40
701f6000 dgram       0      0        0 701fc800        0 700b4dc0
701fc7c0
7017c600 dgram       0      0 14917ed0        0        0        0
/dev/.SRC-unix
/SRCL9gPEe
7017ee00
7007e000 dgram       0      0 1311b458        0        0        0 /tmp/.PMDV1
7007f980
7017c200 stream      0      0 137a6bb0        0        0        0
/tmp/.X11-unix
/X0
7017edc0

.....
.....    netstat -sr
.....

routing:
        0 bad routing redirects
        0 dynamically created routes
        0 new gateways due to redirects
        0 destinations found unreachable
        0 uses of a wildcard route

.....
.....    netstat -nr
.....

Routing tables
Destination      Gateway           Flags   Refs     Use  If   PMTU  Exp  Groups

Route tree for Protocol Family 2 (Internet):
default          9.3.4.41          UGc       0        0  en1     -   -
9.3.4/23         9.3.4.34          U        33   147110  en1     -   -
9.3.187.168      9.3.4.41          UGHW      1     4718  en1     -   -
127/8            127.0.0.1         U         2      775  lo0     -   -

Route tree for Protocol Family 24 (Internet v6):
::1              ::1               UH        0        0  lo0 16896   -

.....
.....    no -a
```

```
.....

         extendednetstats = 0
                  thewall = 524208
               sockthresh = 85
                   sb_max = 1048576
                somaxconn = 1024
       clean_partial_conns = 0
         net_malloc_police = 0
                   rto_low = 1
                  rto_high = 64
                 rto_limit = 7
                rto_length = 13
           inet_stack_size = 16
               arptab_bsiz = 7
                 arptab_nb = 25
                tcp_ndebug = 100
                    ifsize = 8
                  arpqsize = 1
                  ndpqsize = 50
              route_expire = 1
        send_file_duration = 300
                  fasttimo = 200
           routerevalidate = 0
                 nbc_limit = 393136
             nbc_max_cache = 131072
             nbc_min_cache = 1
                  nbc_pseg = 0
              nbc_pseg_limit = 524208
                  strmsgsz = 0
                  strctlsz = 1024
                  nstrpush = 8
                 strthresh = 85
                 psetimers = 20
               psebufcalls = 20
                strturncnt = 15
              pseintrstack = 12288
                 lowthresh = 90
                 medthresh = 95
                  psecache = 1
           subnetsarelocal = 1
                    maxttl = 255
                 ipfragttl = 60
           ipsendredirects = 1
              ipforwarding = 0
                   udp_ttl = 30
                   tcp_ttl = 60
                arpt_killc = 20
             tcp_sendspace = 16384
```

```
          tcp_recvspace = 16384
          udp_sendspace = 9216
          udp_recvspace = 41920
        rfc1122addrchk = 0
        nonlocsrcroute = 0
          tcp_keepintvl = 150
```

Sample output of the nimadm command

In Example B-3, we show you a sample output that you will get when executing the **nimadm** command. This shows the different phases that the operating system performs doing an alternate disk migration.

Example: B-3 Sample output from the alternate disk migration process

```
Initializing the NIM master.
Initializing NIM client srvr80j.itsc.austin.ibm.com.
Verifying alt_disk_migration eligibility.
Initializing log: /var/adm/ras/alt_mig/srvr80j_alt_mig.log
Starting Alternate Disk Migration.

+-----------------------------------------------------------------------------+
Executing nimadm phase 1.
+-----------------------------------------------------------------------------+
Cloning altinst_rootvg on client, Phase 1.
Client alt_disk_install command: alt_disk_install -M 5.2 -C -P1 hdisk4 hdisk5
Calling mkszfile to create new /image.data file.
Checking disk sizes.
Creating cloned rootvg volume group and associated logical volumes.
Creating logical volume alt_hd5.
Creating logical volume alt_hd6.
Creating logical volume alt_hd8.
Creating logical volume alt_hd4.
Creating logical volume alt_hd2.
Creating logical volume alt_hd9var.
Creating logical volume alt_hd3.
Creating logical volume alt_hd1.
Creating logical volume alt_paging00.
Creating logical volume alt_hd10opt.
Creating /alt_inst/ file system.
Creating /alt_inst/home file system.
Creating /alt_inst/opt file system.
Creating /alt_inst/tmp file system.
Creating /alt_inst/usr file system.
Creating /alt_inst/var file system.
Generating a list of files
for backup and restore into the alternate file system...
Backing-up the rootvg files and restoring them to the alternate file system...
```

```
Phase 1 complete.

+-----------------------------------------------------------------------------+
Executing nimadm phase 2.
+-----------------------------------------------------------------------------+
Exporting alt_inst filesystems from client srvr80j.itsc.austin.ibm.com
to NIM master srvr80z.itsc.austin.ibm.com:
Exporting /alt_inst from client.
Exporting /alt_inst/home from client.
Exporting /alt_inst/opt from client.
Exporting /alt_inst/tmp from client.
Exporting /alt_inst/usr from client.
Exporting /alt_inst/var from client.

+-----------------------------------------------------------------------------+
Executing nimadm phase 3.
+-----------------------------------------------------------------------------+
NFS mounting client's alt_inst filesystems on the NIM master:
Mounting srvr80j.itsc.austin.ibm.com:/alt_inst.
Mounting srvr80j.itsc.austin.ibm.com:/alt_inst/home.
Mounting srvr80j.itsc.austin.ibm.com:/alt_inst/opt.
Mounting srvr80j.itsc.austin.ibm.com:/alt_inst/tmp.
Mounting srvr80j.itsc.austin.ibm.com:/alt_inst/usr.
Mounting srvr80j.itsc.austin.ibm.com:/alt_inst/var.

+-----------------------------------------------------------------------------+
Executing nimadm phase 4.
+-----------------------------------------------------------------------------+
nimadm: There is no user customization script specified for this phase.

+-----------------------------------------------------------------------------+
Executing nimadm phase 5.
+-----------------------------------------------------------------------------+
Saving system configuration files.
Expanding /alt_inst/tmp client filesystem.
Filesystem size changed to 98304
Checking for initial required migration space.
Expanding /alt_inst/ client filesystem.
Filesystem size changed to 98304
Expanding /alt_inst/var client filesystem.
Filesystem size changed to 65536
Setting up for base operating system restore.
Restoring base operating system.
Restoring device ODM database.
Merging system configuration files.
Running migration merge method: ODM_merge SWservAt.
Running migration merge method: convert_errnotify.
Running migration merge method: passwd_mig.
Running migration merge method: login_mrg.
```

```
Running migration merge method: user_mrg.
Running migration merge method: secur_mig.
Running migration merge method: mkusr_mig.
Running migration merge method: group_mig.
Running migration merge method: ldapcfg_mig.
Running migration merge method: convert_errlog.
Running migration merge method: merge_smit_db.
Running migration merge method: ODM_merge fix.
Running migration merge method: merge_swvpds.

+-----------------------------------------------------------------------------+
Executing nimadm phase 6.
+-----------------------------------------------------------------------------+
Installing and migrating software.
cat: 0652-050 Cannot open
/srvr80j_alt/alt_inst/usr/sys/inst.data/sys_bundles/GOS.autoi.
Checking space requirements for installp install.
Expanding /alt_inst/usr client filesystem.
Filesystem size changed to 1605632
Installing software with the installp installer.
**************************************************
The installp-output was deleted by the author.
It is included in the logfile anyway.
**************************************************

+-----------------------------------------------------------------------------+
Executing nimadm phase 7.
+-----------------------------------------------------------------------------+
nimadm: There is no user customization script specified for this phase.

+-----------------------------------------------------------------------------+
Executing nimadm phase 8.
+-----------------------------------------------------------------------------+
Creating client boot image.
bosboot: Boot image is 16755 512 byte blocks.
Writing boot image to client's alternate boot disk hdisk4.

+-----------------------------------------------------------------------------+
Executing nimadm phase 9.
+-----------------------------------------------------------------------------+
Unmounting client NFS mounts on the NIM master:
forced unmount of /srvr80j_alt/alt_inst/var
forced unmount of /srvr80j_alt/alt_inst/usr
forced unmount of /srvr80j_alt/alt_inst/tmp
forced unmount of /srvr80j_alt/alt_inst/opt
forced unmount of /srvr80j_alt/alt_inst/home
forced unmount of /srvr80j_alt/alt_inst

+-----------------------------------------------------------------------------+
```

```
Executing nimadm phase 10.
+-------------------------------------------------------------------------------+
Unexporting alt_inst filesystems on client srvr80j.itsc.austin.ibm.com:
exportfs: 1831-184 unexported /alt_inst
exportfs: 1831-184 unexported /alt_inst/home
exportfs: 1831-184 unexported /alt_inst/opt
exportfs: 1831-184 unexported /alt_inst/tmp
exportfs: 1831-184 unexported /alt_inst/usr
exportfs: 1831-184 unexported /alt_inst/var

+-------------------------------------------------------------------------------+
Executing nimadm phase 11.
+-------------------------------------------------------------------------------+
Cloning altinst_rootvg on client, Phase 3.
Client alt_disk_install command: alt_disk_install -M 5.2 -C -P3 hdisk4 hdisk5
## Phase 3 ###################
Verifying altinst_rootvg...
Modifying ODM on cloned disk.
forced unmount of /alt_inst/var
forced unmount of /alt_inst/usr
forced unmount of /alt_inst/tmp
forced unmount of /alt_inst/opt
forced unmount of /alt_inst/home
forced unmount of /alt_inst
forced unmount of /alt_inst
Changing logical volume names in volume group descriptor area.
Fixing LV control blocks...
Fixing file system superblocks...
Bootlist is set to the boot disk: hdisk4

+-------------------------------------------------------------------------------+
Executing nimadm phase 12.
+-------------------------------------------------------------------------------+
Cleaning up alt_disk_migration on the NIM master.
Cleaning up alt_disk_migration on client srvr80j.
```

Abbreviations and acronyms

ACL	Access Control List
AIX	Advanced Interactive Executive
APAR	Authorized Program Analysis Report
API	Application Programming Interface
ATM	Asynchronous Transfer Mode
AuditRM	Audit Log Resource Manager
BLV	Boot Logical Volume
BOS	Base Operating System
CAPP	Controlled Access Protection Profile
CD	Compact Disk
CDE	Common Desktop Environment
CD-R	CD Recordable
CD-ROM	Compact Disk-Read Only Memory
CHRP	Common Hardware Reference Platform
CMP	Certificate Management Protocol
CPU	Central Processing Unit
CSM	Cluster Systems Management
CUoD	Capacity Upgrade on Demand
DASD	Direct Access Storage Device
DB2 UDB	DB2 Universal Database™
DBA	Database Administration
DCE	Distributed Computing Environment
DCUoD	Dynamic Capacity Upgrade on Demand
DFP	Dynamic Feedback Protocol
DGD	Dead Gateway Detection
DLPAR	Dynamic LPAR
DNS	Domain Naming System
DR	Dynamic Reconfiguration
DVD	Digital Video Disk
EAL4+	Evaluation Assurance Level 4+
ERRM	Event Response Resource Manager
ESSL	Engineering and Scientific Subroutine Library
EZNIM	Easy NIM Tool
FAQ	Frequently Asked Questions
FDDI	Fiber Distributed Data Interface
FSRM	File System resource manager
GPFS	General Parallel File System
GUI	Graphical User Interface
HACMP	High Availability Cluster Multi-Processing
HMC	Hardware Management Console
HostRM	Host Resource Manager
HPC	High Performance Computing
HSM	Hierarchical Storage Management
HTTP	Hypertext Transfer Protocol
IBM	International Business Machines Corporation
IP	Internetwork Protocol
IPL	Initial Program Load

ISA Industry Standard Architecture, Instruction Set Architecture

ISICC IBM SAP International Competence Center

ISMP InstallShield Multi-Platform

ITSO International Technical Support Organization

JCE Java Cryptography Extension

JFS Journaled File System

JFS2 Enhanced Journaled File System

JRE Java Runtime Environment

JSSE Java Secure Sockets Extension

LAN Local Area Network

LDAP Lightweight Directory Interchange Format

LPAR Logical Partitioning

LTG Logical Track Group

LVM Logical Volume Manager

MAC Machine Address Code

MCA Micro Channel Architecture

MCM Multichip Module

MPIO Multipath I/O

NBC Network Buffer Cache

NFS Network File System

NIM Network Installation Management

NIS Network Information Service

OEM Original Equipment Manufacturer

ODM Object Data Manager

OS Operating System

PAM Pluggable Authentication Module Support

PCI Peripheral Component Interconnect

PKCS Public-Key Crytography Standards

PKI Public Key Infrastructure

POWER Performance Optimization with Enhanced Risc (Architecture)

PP Physical Partition

PReP POWERPC Reference Platform

PRNG Pseudo-Random Number Generator

PSSP Parallel System Support Program

PTF Program Temporary Fix

PTX Performance Toolbox

P2RSC Power2 Single Chip

QoS Quality of Service

RAM Random Access Memory

RAS Reliability, Availability, and Serviceability

RFC Request for Comments

RIAD Redundant Array of Independent Disks

RMC Resource Monitoring and Control

RSC Power Single Chip

ROM Read Only Memory

RPM Red Hat Package Manager

RSCT Reliable Scalable Cluster Technology

rsh Remote Shell

RTE Run Time Environment

SAP Systeme Anwendungen und Produkte

SCSI Small Computer System Interface

SE Single End

SMIT System Management Interface Tool

SMP	Symmetric Multiprocessor
SMS	System Management Services
SP	IBM RS/6000 Scalable POWER Parallel Systems
SPOT	Shared Product Object Tree
SWVPD	Software Vital Product Database
TCB	Trusted Computing Base
TCP/IP	Transmission Control Protocol/Internet Protocol
TSM	Tivoli Storage Manager
UDF	Universal Disk Format
UDI	Uniform Device Interface
VGDA	Volume Group Descriptor Areas
VIPA	Virtual IP Address
VLAN	Virtual Local Area Network
WAN	Wide Area Network
WLM	Workload Manager
WLRM	Workload Manager Resource Manager
WSM	Web-based System Manager
XLF	XL Fortran

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

IBM Redbooks

For information on ordering these publications, see “How to get IBM Redbooks” on page 236. Note that some of the documents referenced here may be available in softcopy only.

- *AIX 5L Differences Guide Version 5.2 Edition,* SG24-5765
- *AIX 5L Workload Manager (WLM)*, SG24-5977
- *The Complete Partitioning Guide for IBM @server pSeries Servers*, SG24-7039
- *A Holistic Approach to a Reliable Infrastructure for SAP R/3 on AIX*, SG24-5050
- *IBM @server pSeries 670 and pSeries 690 System Handbook*, SG24-7040
- *An Introduction to CSM 1.3 for AIX 5L*, SG24-6859
- *Linux Applications on pSeries*, SG24-6033
- *A Practical Guide for Resource Monitoring and Control (RMC)*, SG24-6615
- *NIM: From A to Z in AIX 4.3*, SG24-5524

Other publications

These publications are also relevant as further information sources:

- *AIX Version 4.3 System Management Guide: Operating System and Devices*, SC23-4126
- *AIX 5L for POWER Version 5.1 Release Notes*, GI10-0729
- *AIX 5L Version 5.1 Installation Guide and Reference*, SC23-4374
- *AIX 5L Version 5.1 Network Installation Management Guide and Reference*, SC23-4385
- *AIX 5L Version 5.2 Commands Reference, Volume 1 to Volume 6*, found at:

  `http://publib16.boulder.ibm.com/cgi-bin/ds_rslt#1`

- *AIX 5L Version 5.2 Installation Guide and Reference*, SC23-4389
- *AIX 5L Version 5.2 Performance Management Guide*, found at:

  http://publib16.boulder.ibm.com/doc_link/en_US/a_doc_lib/aixbman/prftungd/prftungdtfrm.htm
- *AIX 5L Version 5.2 Release Notes*, GI10-0739
- *AIX 5L Version 5.2 Security Guide*, found at:

  http://publib16.boulder.ibm.com/pseries/en_US/aixbman/security/securitytfrm.htm
- *AIX 5L Version 5.2 Web-based System Manager Administrator Guide*, found at:

  http://publib16.boulder.ibm.com/doc_link/en_US/a_doc_lib/aixbman/wsmadmn/wsmadmntfrm.htm
- *Hardware Management Console Installation and Operations Guide*, SA38-0590
- *IBM SAP Technical Brief Migration of Oracle 8.1.7/AIX 4.3.3 SAP installation To Oracle 9.2/AIX 5.1 SAP installation* (Version 1.00, December 2002)
- SAP Guide *Upgrade to Oracle Version 9.2.0: UNIX*

Online resources

These Web sites and URLs are also relevant as further information sources:

- AIX 5.2 Documentation in Information Center

  http://publib16.boulder.ibm.com/pseries/en_US/infocenter/base/aix52.htm
- AIX Documentation

  http://publib16.boulder.ibm.com/cgi-bin/ds_form?lang=en_US&viewset=AIX
- AIX Software

  http://www.ibm.com/servers/aix/products/
- AIX Toolbox for Linux Applications

  http://www.ibm.com/servers/aix/products/aixos/linux/
- AIX Version 4.3 books

  http://publib.boulder.ibm.com/cgi-bin/ds_form?lang=en_US

- Counterpane Labs: Yarrow

  http://www.counterpane.com/yarrow.html

- Current Release Notes for AIX

  http://publib.boulder.ibm.com/pseries/aixgen/relnotes/current_relnotes.html

- DB2 APARs

  http://www.ibm.com/cgi-bin/db2www/data/db2/udb/winos2unix/support/hiperapar.d2w/report

- DB2 End-of-Service

  http://www.ibm.com/cgi-bin/db2www/data/db2/udb/winos2unix/support/newsletter.d2w/n20020101

- Fix Delivery Center for AIX Version 5

  http://techsupport.services.ibm.com/server/aix.fdc

- IBM @server pSeries Information Center

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

- IBM Tivoli Storage Manager

  http://publib.boulder.ibm.com/tividd/td/IBMStorageManagerforAIX5.1.html

- Kornshell home page

  http://www.kornshell.com

- OpenSSH for AIX

  http://www.ibm.com/servers/esdd/articles/openssh_updated.html

- pSeries and RS/6000 Microcode Updates

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

- Release Notes Index for AIX 5.1 and Related Products

  http://publib.boulder.ibm.com/pseries/aixgen/relnotes/51relnotes.current.html

- Release Notes Index for AIX 5.2 and related products

  http://publib.boulder.ibm.com/pseries/aixgen/relnotes/52relnotes.current.html

- SAP Service Marketplace (requires registration)

  http://service.sap.com/instguides
  http://service.sap.com/notes
  http://service.sap.com/systemmanagement

- Search and download selective fixes for AIX 5.1

  http://techsupport.services.ibm.com/server/aix.fixsearch51

- Security Guide: OpenSSH Software Tools

  http://publib16.boulder.ibm.com/pseries/en_US/aixbman/security/openssh.htm

How to get IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this Web site:

ibm.com/redbooks

Index

B

C

D

E

F

G

H

N

O

P

Q

R

S

T

U

V

W

X

Z

IBM Redbooks

AIX Version 4.3 to 5L Migration Guide

(0.2”spine)
0.17”<->0.473”
90<->249 pages

AIX Version 4.3 to 5L Migration Guide

Latest enhancements in AIX 5L Version 5.1 and Version 5.2

Planning and preparation hints and tips for your migration

Step-by-step migration tasks by media and NIM

This IBM Redbook will provide AIX system administrators with the key information they need to know in order to make a safe migration to AIX 5L. It covers the steps that are required when planning, preparing, and actually migrating a system to the latest available level of AIX (currently AIX 5L Version 5.2) from AIX Version 4.3.

Before discussing the actual migration steps, we first point out the latest enhancements and new functions that have been delivered with AIX 5L Version 5.1 and Version 5.2 and why we recommend that you upgrade your system.

We then explain how to plan, prepare, and perform the migration, focusing on providing hints, tips, and sources of additional information. We will explain detailed step-by-step actual migration by media and Network Installation Management (NIM).

This redbook also covers post-migration issues, including NIM master and alternate disk migration specific tasks. Lastly, we will summarize major applications issues when you migrate your system from AIX Version 4.3 to AIX 5L.

This redbook will help you design your solution to migrate your AIX system to the latest available version of AIX.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-6924-00 ISBN 073842949X