International Technical Support Organization

SG24-4653-00

A Holistic Approach to AIX V4.1 Migration, Volume 2 TCP/IP, SNA, HACMP and Multiple Systems

May 1996

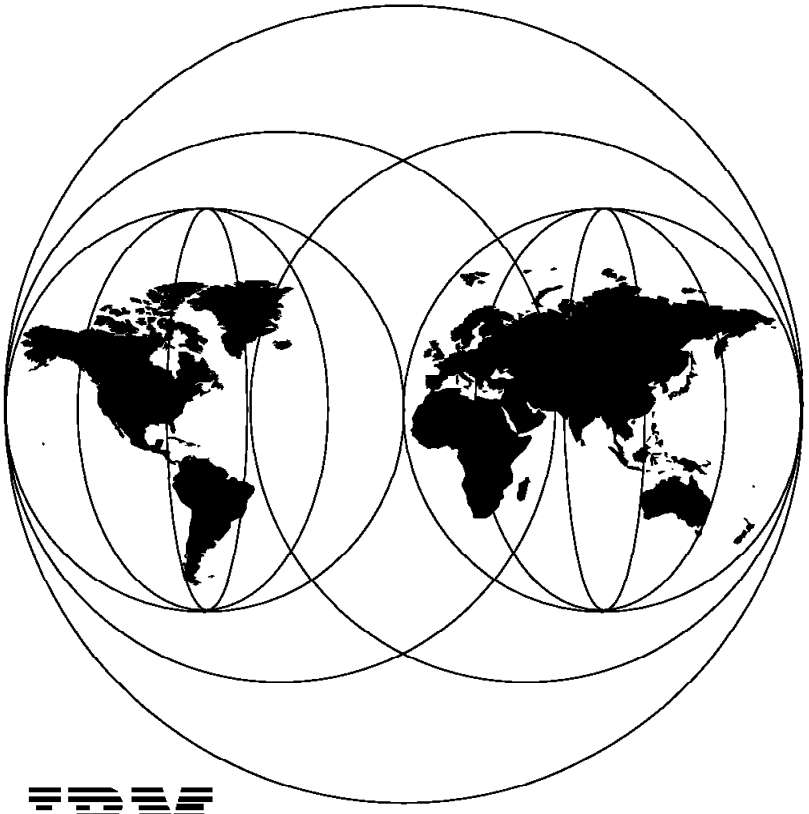

IBM

International Technical Support Organization

Austin Center

International Technical Support Organization

SG24-4653-00

A Holistic Approach to AIX V4.1 Migration, Volume 2 TCP/IP, SNA, HACMP and Multiple Systems

May 1996

Take Note!

Before using this information and the product it supports, be sure to read the general information under “Special Notices” on page xv.

First Edition (May 1996)

This edition applies to Version 3, Release 2, Modification Level 5 of the AIX Operating System, Program Number 5765-030, Version 4, Release 1, Modification Levels 3 and 4 of the AIX Operating System, Program Number 5765-393, and associated IBM and Independent Software Vendor Program Products.

Order publications through your IBM representative or the IBM branch office serving your locality. Publications are not stocked at the address given below.

An ITSO Technical Bulletin Evaluation Form for reader's feedback appears facing Chapter 1. If the form has been removed, comments may be addressed to:

IBM Corporation, International Technical Support Organization
Dept. JN9B Building 45 Internal Zip 2834
11400 Burnet Road
Austin, Texas 78758-3493

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Abstract

A Holistic Approach to AIX V4.1 Migration, Volume 2, is the final part of the Holistic Migration trilogy:

- A Holistic Approach to AIX V4.1 Migration, Planning Guide
- A Holistic Approach to AIX V4.1 Migration, Volume 1
- A Holistic Approach to AIX V4.1 Migration, Volume 2

This series takes a broad-based approach to the migration of your RISC System/6000 computing environment from AIX Version 3.2 to AIX Version 4.1. It covers the migration of complete systems—systems that include not only the AIX operating system but also applications, networking customizations, and the many other factors that make each installation unique.

This volume deals specifically with the migration of systems that utilize communications over TCP/IP and SNA networks and shows system administrators the quickest methods for upgrading and reconfiguring their system.

It also details the procedures for using Network Installation Management, a new feature of AIX Version 4.1, to perform unattended migrations of multiple systems connected on a local area network.

Finally, it presents the steps required to migrate a node in a High Availability Cluster Multi-Processing (HACMP) environment with a minimum of risk and downtime, and a discussion of the compatibility between HACMP versions for heterogeneous operations is provided.

The holistic migration trilogy guides the system administrator past the pitfalls and risks associated with any major system maintenance and ensures the smoothest possible upgrade experience.

This document is written for current users of AIX Version 3.2.5 and associated products. Basic knowledge of the AIX operating system and the specific environment to be migrated is assumed.

(296 pages)

Contents

Figures

Tables

Special Notices

This publication is intended to help administrators of systems that are currently running AIX Version 3.2 to migrate their systems to AIX Version 4.1. The information in this publication is not intended as the specification of any programming interfaces that are provided by AIX. See the PUBLICATIONS section of the IBM Programming Announcement for AIX for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM′s product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM′s intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter 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 the IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, NY 10594 USA.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM (VENDOR) products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer′s ability to evaluate and integrate them into the customer′s operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

The following document contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples contain 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.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

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

ACF/VTAM	Advanced Peer-to-Peer Networking
AIX	AIX/6000
AIXwindows	AnyNet
APPN	C Set ++
CICS	CICS/6000
ESCON	HACMP/6000
IBM	InfoExplorer
MERVA	NetView
Portmaster	PowerPC
PowerPC 601	PowerPC 604
PS/2	RISC System/6000
RS/6000	S/370
S/390	SP
SystemView	VTAM
Xstation Manager	

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc.

DynaText is a trademark of Electronic Book Technologies, Inc.

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Microsoft, Windows, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

DCE, Motif	Open Software Foundation
DynaText	Electronic Book Technologies, Inc.
NFS	Sun Microsystems, Inc.
SCSI	Security Control Systems, Inc.
POSIX	Institute of Electrical and Electronic Engineers

Other trademarks are trademarks of their respective companies.

Preface

The Webster's Ninth New Collegiate Dictionary defines holistic as:

1: *of or relating to holism* ***2:*** *relating to or concerned with wholes or with complete systems rather than with the analysis of, treatment of, or dissection into parts.*

This second definition describes the approach we have taken with the migration of systems from AIX Version 3.2 to Version 4.1. In the real world, a migration doesn't involve just the operating system. A computing environment includes applications from both IBM and independent software vendors and a great deal of site-specific customization of system and user environments. In this case, there are always going to be unique migration actions that cannot be performed automatically by migration tools.

Our testing included systems running applications such as HACMP and the Oracle database. It also included systems and groups of systems with extensive customizations in TCP/IP and SNA communications protocols. While it is obviously not possible to cover every specific environment, we hope that the environments we have chosen will provide a wide enough cross section to give system administrators clues or hints as to the problems they may discover in their own environments.

This document is the final part of a three-volume series on migration. The contents of each volume are listed below:

- *A Holistic Approach to AIX V4.1 Migration, Planning Guide*

  This covers the advance planning required for a successful migration. It includes such topics as documenting the existing system, determining the required levels of software for the new environment, verifying that hardware is supported at the new level, and scheduling the migration.

- *A Holistic Approach to AIX V4.1 Migration, Volume 1*

  This is the first of two volumes to assist system administrators in actually performing the migration. It includes information on the base operating system, migrating from Uni-Processor systems to those using Symetric Multi-Processors, and the migration of systems using the Oracle database. It also includes a detailed description of the migration process.

- *A Holistic Approach to AIX V4.1 Migration, Volume 2*

  The final volume covers communications issues, such as the migration of systems with complex network configurations using TCP/IP and SNA. It also includes sections on the migration of systems running HACMP, and the migration of large numbers of systems in LAN environments.

How This Document is Organized

The document is organized as follows:

- Part 1, "TCP/IP"

  This part describes the migration of systems configured for TCP/IP communications. It considers such TCP/IP facilities as Domain Name Server, Network File System, and Network Information Services.

- Part 2, “SNA”

  These chapters describe the migration of systems running either SNA Services/6000 Version 1.2 or SNA Server/6000 Version 2.1 to AIX Communications Server Version 4. It includes the migration of SNA Transaction Programs, and the use of X.25 and System/370 channel links.

- Part 3, “Migration of Multiple Systems”

  This part covers the migration of multiple systems—in LAN environments, using the AIX Network Installation Management function.

- Part 4, “High Availability Cluster Multi-Processing”

  In a High Availability Cluster Multi-Processing environment, it is important to achieve the migration of the cluster nodes with a minimum of downtime and risk. This part describes the potential problems with migration of cluster nodes, and how to avoid these problems.

Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this document.

- *AIX Version 4.1 Getting Started*, SC23-2527
- *AIX Version 4.1 Installation Guide*, SC23-2550
- *Common Diagnostics Information Manual*, SA23-2765
- *AIX Version 3.2 Commands Reference*, GBOF-1802
- *AIX Version 3.2 System Management Guide: Operating System and Devices*, GC23-2486
- *AIX Version 4.1 Commands Reference*, SBOF-1851
- *AIX Version 4.1 System Management Guide: Operating System and Devices*, SC23-2525
- *AIX Version 4.1 Network Installation Management Guide and Reference*, SC23-2627
- *AIX Version 4.1 Files Reference*, SC23-2512
- *All About AIX Version 4.1*

  This book was never officially published in hardcopy; however, it is available on the World-Wide Web at the URL:

  `http://www.developer.ibm.com/sdp/library/ref/about4.1/df4main.html`

  Within IBM, *All About AIX Version 4.1* is also available as ABOUT4_1 PACKAGE on the MKTTOOLS disk.

International Technical Support Organization Publications

A complete list of International Technical Support Organization redbooks, with a brief description of each, may be found in:

International Technical Support Organization Bibliography of Redbooks, GG24-3070.

How Customers Can Get Redbooks and Other ITSO Deliverables

Customers may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **IBMLINK**

  Registered customers have access to PUBORDER to order hardcopy, to REDBOOKS disk to obtain BookManager BOOKs

- **IBM Bookshop** — send orders to:

  usib6fpl@ibmmail.com (United States)
  bookshop@dk.ibm.com (Outside United States)

- **Telephone orders**

  1-800-879-2755 (United States)
  354-9408 (Australia)
  359-2-731076 (Bulgaria)
  42-2-67106-250 (Czech Republic)
  593-2-5651-00 (Ecuador)
  03-69-78901 (Israel)
  905-627-1163 (Mexico)
  064-4-57659-36 (New Zealand)
  027-011-320-9299 (South Africa)
  01256-478166 (United Kingdom)
  32-2-225-3738 (Belgium)
  1-800-IBM-CALL (Canada)
  45-934545 (Denmark)
  01805-5090 (Germany)
  0462-73-6669 (Japan)
  31-20513-5100 (The Netherlands)
  507-639977 (Panama)

- **Mail Orders** — send orders to:

  IBM Publications
  P.O. Box 9046
  Boulder, CO 80301-9191
  USA

  IBM Direct Services
  Sortemosevej 21
  DK-3450 Allerød
  Denmark

- **Fax** — send orders to:

  1-800-445-9269 (United States)
  32-2-225-3478 (Belgium)
  905-316-7210 (Canada)
  593-2-5651-45 (Ecuador)
  03-69-59985 (Israel)
  31-20513-3296 (The Netherlands)
  507-693604 (Panama)
  01256-843173 (United Kingdom)
  359-2-730235 (Bulgaria)
  42-2-67106-402 (Czech Republic)
  07032-15-3300 (Germany)
  0462-73-7313 (Japan)
  064-4-57659-16 (New Zealand)
  027-011-320-9113 (South Africa)

- **1-800-IBM-4FAX (United States only)** — ask for:

  Index # 4421 Abstracts of new redbooks
  Index # 4422 IBM redbooks
  Index # 4420 Redbooks for last six months

- **Direct Services**

  Send note to softwareshop@vnet.ibm.com

- **Redbooks Home Page on the World Wide Web**

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

- **E-mail (Internet)**

  Send note to redbook@vnet.ibm.com

- **Internet Listserver**

With an Internet E-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an E-mail note to announce@webster.ibmlink.ibm.com with the keyword subscribe in the body of the note (leave the subject line blank). A category form and detailed instructions will be sent to you.

How IBM Employees Can Get Redbooks and ITSO Deliverables

Employees may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **PUBORDER** — to order hardcopies in United States
- **GOPHER link to the Internet**

  Type GOPHER
  Select IBM GOPHER SERVERS
  Select ITSO GOPHER SERVER for Redbooks

- **Tools disks**

  To get LIST3820s of redbooks, type one of the following commands:

  ```
  TOOLS SENDTO EHONE4 TOOLS2 REDPRINT GET GG24xxxx PACKAGE
  TOOLS SENDTO CANVM2 TOOLS REDPRINT GET GG24xxxx PACKAGE (Canadian users only)
  ```

  To get lists of redbooks:

  ```
  TOOLS SENDTO WTSCPOK TOOLS REDBOOKS GET REDBOOKS CATALOG
  TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET ITSOCAT TXT
  TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET LISTSERV PACKAGE
  ```

  To register for information on workshops, residencies, and redbooks:

  ```
  TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ITSOREGI 1996
  ```

  For a list of product area specialists in the ITSO:

  ```
  TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ORGCARD PACKAGE
  ```

- **Redbooks Home Page on the World Wide Web**

  http://w3.itso.ibm.com/redbooks/redbooks.html

  IBM employees may obtain LIST3820s of redbooks from this page.

- **ITSO4USA category on INEWS**
- **IBM Bookshop** — send orders to:

  USIB6FPL at IBMMAIL or DKIBMBSH at IBMMAIL

- **Internet Listserver**

  With an Internet E-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an E-mail note to announce@webster.ibmlink.ibm.com with the keyword subscribe in the body of the note (leave the subject line blank). A category form and detailed instructions will be sent to you.

Acknowledgments

This project was designed and managed by:

Yves Bex
International Technical Support Organization, Austin Center

Cameron Ferstat
International Technical Support Organization, Austin Center

The authors of this document are:

Indulis Bernsteins
IBM Australia

Pascale Delava
IBM Belgium

Lars Ellingsberg
IBM Norway

Andreas Hermelink
IBM Germany

Motonobu Koh
IBM Japan

Zhu Li
IBM China

Yasuhiro Saitah
IBM Japan

Harald Schneider
IBM Germany

This publication is the result of a residency conducted at the International Technical Support Organization, Austin Center.

Thanks to the following people for the invaluable advice and guidance provided in the production of this document:

Rich Avery
IBM Austin

Julie Craft
IBM Austin

John Ellis
IBM Austin

John Kennedy
IBM Austin

Marshall Lamb
IBM Raleigh

Paul Landay
IBM Raleigh

Al Mitchell
International Technical Support Organization, Austin Center

Barry Nusbaum
International Technical Support Organization, Raleigh Center

Marc Stephenson
IBM Austin

David Thiessen
International Technical Support Organization, Austin Center

Marcus Brewer
Technical Editor
International Technical Support Organization, Austin Center

Part 1. TCP/IP

Chapter 1. Migrating TCP/IP

In this chapter, we cover the migration of some TCP/IP services from AIX Version V3.2.5 to AIX Version V4.1.4 using the Migration Install method.

TCP/IP migration is done as part of the Migration Install of AIX Version V4.1.4; so a successful TCP/IP migration relies on a successful AIX migration. You should have already read the AIX migration chapter in *A Holistic Approach to AIX V4.1 Migration, Volume 1*, SG24-4652, which covers the general issues in migrating from AIX Version 3 to AIX Version 4.

Attention!

The Migration Install process for AIX V4.1.4 (and earlier versions) has a problem with the correct migration of NFS and NIS configuration files. You will need to take manual steps to properly complete the migration process (see 1.8.3, “NFS Differences and Migration Experiences” on page 31).

You should also be aware of the following limitations of the migration process:

- Adapters may be renamed during the migration process. The migration process renames adapters in the order in which they are found in the physical adapter slots.
- Volume groups (and associated file systems) may not be automatically made available after a migration. This is important if your TCP/IP configuration files are in a file system on one of these volume groups.

You should establish test procedures so that you can verify if the migration of TCP/IP services was successful. This may mean scheduling time and resources to do this testing. For example, you may need to make sure someone is at a remote site when you test remote printing.

For more information about planning your migration to AIX V4.1, please see *A Holistic Approach to AIX V4.1 Migration, Planning Guide*, SG24-4651.

Experienced TCP/IP system administrators may want to use Table 1 on page 7 for guidance on migrating TCP/IP services, and refer to the text only when necessary.

1.1.1 Related Publications

For more information, please consult:

- *System Management Guide: Communications and Networks*. This is available in InfoExplorer or as hardcopy manual GC23-2487 (AIX V4.1) or SC23-2526 (AIX V3.2).
- *TCP/IP Tutorial and Technical Overview*, GG24-3376
- *DNS and Bind, P.Albitz and C.Liu, O'Reilly 1994*, also available from IBM as publication number SR28-4970
- *Managing NFS and NIS, Hal Stern, O'Reilly 1994*, also available from IBM as publication number SR28-4969

1.2 AIX Version V4.1 TCP/IP Packaging and Features

This section covers the AIX V4.1.4 TCP/IP packaging and features.

1.2.1 AIX V4.1 Client versus Server

With the announcement of AIX V4.1, IBM released a client and server version of AIX. These are two different packages of AIX and associated filesets which are priced and customized for different requirements. In the original client package of AIX V4.1.1, many TCP/IP servers and functions were not included since they were thought to be mainly used on larger server systems. For example, the AIX client package did not include an NFS server or a DNS name server.

In the true spirit of listening to customers, it was quickly realized that many TCP/IP clients were also servers. With the announcement of AIX V4.1.3, the client version of AIX included all of the TCP/IP functionality of the server version. The client version of AIX V4.1.4 also includes all of the server functionality of server version.

Some filesets which use TCP/IP functions but are not part of the TCP/IP filesets (bos.net.*) are still only available in AIX Version 4 for Servers, and are not included with AIX Version 4 for Clients. Examples of this include Network Installation Management (NIM) and Xstation Manager.

1.2.2 TCP/IP Features in AIX V4.1

TCP/IP features in AIX Version 4.1 include:

- **DHCP support:**

  Dynamic Host Configuration Protocol (introduced with AIX V4.1.4) allows system administrators to automatically assign network addresses to a particular hostname. This makes system administration easier, especially for systems that move around on a network, such as laptop PCs.

- **PPP support:**

  Point-to-Point-Protocol (PPP, introduced with AIX V4.1.4) for TCP/IP over asynchronous lines is now supported in addition to SLIP.

- **sliplogin command:**

  The `sliplogin` command simplifies the support of SLIP over dial-up lines and allows dial-up ports to support both normal dialup users and SLIP connections.

- **IP multicast:**

  IP multicasting is used to broadcast to logical groups of network addresses. It is also used for multimedia applications, such as live video over the Internet.

- **Packet drops:**

  The `netstat` command has been enhanced to show the number of packet drops.

- **Promiscuous mode:**

  This mode is supported for the integrated Ethernet, Ethernet and FDDI adapters. It allows the `iptrace` command to be used to look at all packets on the LAN, not just packets destined for adapters on this system.

- **OSPF routing:**

  Open Shortest Path First routing is supported (using gated, based on Cornell gated 3.0.2). This routing algorithm is an improvement over the vector-distance algorithms previously used.

- **Resolver API:**

  New programming interfaces in the TCP/IP library allow programmers to choose if they want BIND (named), NIS or /etc/hosts to resolve hostnames to IP addresses. Programmers can also add other name services.

1.3 Migrating TCP/IP to AIX V4.1.4

This section describes the environment used for our tests and provides a TCP/IP migration summary that can be used directly by experienced TCP/IP administrators.

1.3.1 Migration Environment

Attention!

The Migration Install process for AIX V4.1.4 (and earlier versions) has a problem with the migration of NIS and NFS configuration files. You will need to take manual steps to properly complete the migration process (see 1.8.3, “NFS Differences and Migration Experiences” on page 31).

The migration which was tested was AIX Version V3.2.5 migrating to AIX Version V4.1.4. The migration process tested was Migration Install (selected from the AIX V4.1.4 installation menus). The TCP/IP services which were migrated and tested were:

- Basic networking: basic TCP/IP configuration with default gateway
- Domain Name System (DNS): primary and secondary name servers for a subdomain
- Network File System (NFS): file server and client
- Network Information Service (NIS): master NIS server and NIS client
- r commands
- ftp: including anonymous ftp server

Using the Migration Install method of installing AIX V4.1.4, the migration appeared to proceed smoothly and completed leaving a functioning AIX V4.1.4 system. However, after the end of testing, a problem with the automatic migration process for TCP/IP was discovered.

The result of this problem was that some files from AIX V3.2.5 were left in place instead of being replaced by the default AIX V4.1.4 files. The filesets affected were bos.net.nfs.client and bos.net.nis.client. If you are using NIS or NFS, please refer to 1.8.3, “NFS Differences and Migration Experiences” on page 31 for more information.

If you choose not to use Migration Install and instead choose to use New and Complete Overwrite or Preservation Install, then you should:

- Save all of your TCP/IP configuration files manually, or be sure that you can quickly and reliably recover these files from your system backup. You can use the information in 1.4.1, “Files and Filesets” on page 9 as a guideline for which files to preserve. Note that the way your system is configured may mean there are also other TCP/IP configuration files to backup and recover.
- After migration, reload all of your configuration files onto the system, into a different directory.
- Compare the default AIX V4.1.4 with your own AIX V3.2.5 files, and update the default AIX V4.1.4 files with any local customizations required.
- Reconfigure any devices or other items whose configuration information has been lost. This includes loading any additional software and device drivers that are required.
- Follow the sections in this chapter for the TCP/IP services you are interested in.

1.3.2 TCP/IP Migration Summary

The following table is a summary of the steps that need to be taken for the successful migration of TCP/IP services. This can serve as a quick reference for experienced system administrators.

Table 1 (Page 1 of 2). TCP/IP Migration Summary. AIX V4.1.4 Migration Install

TCP/IP Function	Action required before migration	Action required after migration	Notes
ALL	Refer to *A Holistic Approach to AIX V4.1 Migration, Volume 1* to successfully migrate AIX.	Refer to *A Holistic Approach to AIX V4.1 Migration, Volume 1* for AIX post-migration steps.	AIX must be successfully migrated for TCP/IP migration to succeed.
	Be aware that adapter names may be swapped.	Swap cables or reconfigure adapters in the location to match the original adapter in that location. Adapter configuration information is in /tmp/reconfig* files.	The migration process may result in adapter names being swapped.
	Make a backup copy of /etc/rc.tcpip (say to /etc/rc.tcpip.aix3).	If necessary, manually merge your changes into the new /etc/rc.tcpip.	In some cases, Migration Install of AIX V4.1 does not correctly update /etc/rc.tcpip.
Default gateway	Note the address for the default gateway.	If necessary, use `smit mktcpip` to add in the default gateway.	In some cases, Migration Install of AIX V4.1 does not correctly update /etc/rc.tcpip.
IP forwarding	Note if this is set to 1 (=on) using the `no -o ipforwarding` command.	Use the `no -o ipforwarding=1` command if you need the ability to forward packets not destined for this system	ipforwarding was set to 1 (=on) by default in AIX V3.2.5, but the default is 0 (=off) by default in AIX V4.1.4.
DNS	Request the fix for secondary name server that is not reloading the zone correctly (APAR IX55267). Note if `named` daemon is started automatically.	Get the fix for the reload problem, or use `ln -s /usr/sbin/named-xfer /etc/named-xfer`. If necessary, use `smit stnamed` to select the `named` daemon to be started automatically.	An AIX V4.1.4 secondary name server has a known problem reloading zone data from the primary. In some cases, AIX V4.1.4 /etc/rc.tcpip sets `named` to not start automatically.
r commands	No action required.	No action required.	The r commands `rlogin`, `rsh`, `rexec`, and `rcp` worked as expected.
NFS	No action required.	Manual recovery and merge of `/etc/rc.nfs` and `/etc/nfs.clean` is required. See 1.8.3, “NFS Differences and Migration Experiences” on page 31.	Known problem with NFS migration does not update /etc/rc.nfs and other files. NFS did not have any apparent problems using the AIX V3.2.5 configuration files, but manual recovery/merge is required.
NIS	Disable NIS client daemon `ypbind` in /etc/rc.nfs before migration. Required for NIS clients and servers using `ypbind`.	Manual recovery and merge of `/var/yp/Makefile` and `/var/yp/updaters` is required. See 1.8.3, “NFS Differences and Migration Experiences” on page 31. If required, reenable NIS client daemon `ypbind` after recovery/merge.	Known problem with NIS migration does not update /etc/rc.nfs and other files. AIX V3.2.5 /etc/rc.nfs file will start the NIS client daemon but not the NIS server daemons after AIX V4.1.4 migration.

Table 1 (Page 2 of 2). TCP/IP Migration Summary. AIX V4.1.4 Migration Install

TCP/IP Function	Action required before migration	Action required after migration	Notes

ftp	No action required.	For anonymous ftp server:
		`cp /usr/bin/ls /home/ftp_dir/bin`
		`cp /usr/ccs/lib/libc.a \`
		`/home/ftp_dir/lib/libc.a`
			`ls` and libc.a need to be updated to AIX V4.1.4; otherwise `ls -l` will not work for anonymous ftp clients.

SNMP (not tested)	Note if the `snmpd` daemon is started automatically.	If required, use `chrctcp -S -a snmpd` to set the `snmpd` daemon to be started automatically.	AIX V4.1.4 migration of /etc/rc.tcpip may set `snmpd` to not start automatically.

1.4 How AIX V4.1 Upgrades TCP/IP

If you choose a Migration Install of AIX Version 4, the TCP/IP packages you have installed on AIX V3.2.5 are migrated to their corresponding filesets in AIX V4.1.4. A fileset is the smallest unit of AIX V4.1 software which can be installed on a system. A software package consists of one or more filesets. The migration uses `installp` to do much of its work. The `installp` command has been enhanced in AIX V4.1 to support automatic saving and updating of configuration files. This is used by the TCP/IP filesets to make migration easy and automatic.

A Migration Install only destroys the contents of /tmp and leaves most configuration files intact (or updates them to work with AIX V4.1.4). You should *not* have your TCP/IP configuration files in /tmp.

In most cases, it is a fairly simple matter to migrate TCP/IP services from AIX Version V3.2 to AIX Version V4.1.

1.4.1 Files and Filesets

During migration, the list of AIX V3.2.5 installed software is checked, and the appropriate replacement AIX V4.1.4 TCP/IP filesets are installed.

Each fileset has a list of configuration files which need to be processed during migration. These files are either left alone, replaced by an AIX V4.1.4 default file (with the original file saved), or automatically processed to make the file compatible with the new AIX V4.1.4 format. The action to be taken is controlled by a keyword next to the filename. TCP/IP uses the standard AIX V4.1 installation methodology, which allows filesets to perform pre-installation and post-installation processing. For more information, see the *How It Works* chapter in *A Holistic Approach to AIX V4.1 Migration, Volume 1*.

Keywords for configuration files:

preserve replaces the AIX V4.1.4 default file with the original AIX V3.2.5 file which was saved in /tmp/bos or /lpp/save.config.

user_merge installs the AIX V4.1.4 default file on the system, and leaves a copy of the original AIX V3.2.5 file in /tmp/bos or /lpp/save.config. After Migration Install, you must manually merge changes from the original file with the AIX V4.1.4 file.

auto_merge automatically updates the original file to be compatible with AIX V4.1.4 during the installation of the fileset.

hold_new replaces the AIX V4.1.4 default file with the original AIX V3.2.5 file which was saved in /tmp/bos or /lpp/save.config. The AIX V4.1.4 default file is saved in /tmp/bos or /lpp/save.config for future reference.

other handled the same way as user_merge.

Following are lists of files which are processed by the installation of TCP/IP filesets during a Migration Install of AIX V4.1.4. The files are listed according to the TCP/IP fileset that is installed.

Attention!

You should note the files marked **user_merge** and **other** in the following lists, and check that any customization done in these files on your AIX V3.2.5 system is also in these files on your migrated AIX V4.1.4 system.

The files which are processed as **user_merge** and **other** are shown in **bold** to assist you.

1.4.1.1 Files Processed by TCP/IP Client Installation

Files affected by the TCP/IP client fileset installation (bos.net.tcp.client) during migration to AIX V4.1.4 are:

```
/etc/3270.keys preserve
/etc/3270keys.hft preserve
/etc/3270_arab_kyb.map preserve
/etc/aliases preserve
/etc/bootptab preserve
/etc/dhcpcd.ini preserve
/etc/hosts preserve
/etc/hosts.equiv preserve
/etc/hosts.lpd preserve
/etc/inetd.conf preserve
/etc/map3270 preserve
/etc/mib.defs user_merge
/etc/protocols preserve
/etc/rc.bsdnet hold_new
/etc/rc.net hold_new
/etc/rc.net.serial preserve
/etc/rc.tcpip hold_new
/etc/rpc preserve
/etc/sendmail.cf preserve
/etc/sendmail.nl preserve
/etc/services preserve
/etc/syslog.conf preserve
/etc/slip.hosts preserve
/etc/slip.login preserve
/etc/slip.logout preserve
/etc/snmpd.conf hold_new
/etc/snmpd.peers preserve
/etc/tcp.clean preserve
/etc/telnet.conf preserve
/usr/lib/smdemon.cleanu preserve
/usr/samples/tcpip/README user_merge
/usr/samples/tcpip/anon.ftp preserve
/usr/samples/tcpip/netrc preserve
/usr/samples/tcpip/networks preserve
/usr/samples/tcpip/resolv.conf preserve
/usr/samples/tcpip/rhosts preserve
/usr/samples/tcpip/tftpaccess.ctl preserve
/usr/sbin/chservices user_merge
/usr/sbin/chsubserver user_merge
/usr/sbin/mktcpip user_merge
/usr/sbin/slipcall preserve
/usr/share/lib/Mail.rc preserve
```

1.4.1.2 Files Processed by TCP/IP Server Installation

Files affected by the TCP/IP server fileset installation (bos.net.tcp.server) during migration to AIX V4.1.4 are:

```
/etc/dhcprd.cnf preserve
/etc/dhcpsd.cnf preserve
/etc/gated.conf other
/etc/securetcpip preserve
/usr/samples/snmpd/ethernet.my user_merge
/usr/samples/snmpd/fddi.my user_merge
/usr/samples/snmpd/generic.my user_merge
/usr/samples/snmpd/ibm.my user_merge
/usr/samples/snmpd/mibII.my user_merge
/usr/samples/snmpd/smi.my user_merge
/usr/samples/snmpd/token_ring.my user_merge
/usr/samples/snmpd/unix.my user_merge
/usr/samples/snmpd/view.my user_merge
```

1.4.1.3 Files Processed by NFS Client Installation

Attention!

Note that the following files are not processed correctly during a Migration Install of AIX V4.1.4. See 1.8.3, “NFS Differences and Migration Experiences” on page 31.

Files affected by the NFS client fileset installation (bos.net.nfs.client) during migration to AIX V4.1.4 are:

```
/etc/nfs.clean     user_merge
/etc/rc.nfs        user_merge
```

1.4.1.4 Files Processed by NIS Client Installation

Attention!

Note that the following files are not processed correctly during a Migration Install of AIX V4.1.4. See 1.8.3, “NFS Differences and Migration Experiences” on page 31

Files affected by the NIS client fileset installation (bos.net.nis.client) during migration to AIX V4.1.4 are:

```
/var/yp/Makefile   user_merge
/var/yp/updaters   user_merge
```

1.4.1.5 Files Processed by uucp Installation

Note that uucp migration was not tested. Files affected by the uucp fileset installation (bos.net.uucp) during migration to AIX V4.1.4 are:

```
/etc/uucp/Dialcodes preserve
/etc/uucp/Maxuuscheds preserve
/etc/uucp/Maxuuxqts preserve
/etc/uucp/Permissions preserve
/etc/uucp/Poll preserve
/etc/uucp/Systems preserve
/usr/sbin/uucp/remote.unknown preserve
/usr/sbin/uucp/uudemon.admin preserve
/usr/sbin/uucp/uudemon.cleanu preserve
```

```
/usr/sbin/uucp/uudemon.hour preserve
/usr/sbin/uucp/uudemon.poll preserve
```

1.4.1.6 Files Processed by Asynchronous Terminal Emulation (ATE) Installation

Note that ATE migration was not tested. Files affected by the ATE fileset installation (bos.net.ate) during migration to AIX V4.1.4 are:

```
/usr/lib/dir preserve
```

1.4.1.7 Files Processed by NCS Installation

Note that NCS migration was not tested. Files affected by the NCS fileset installation (bos.net.ncs) during migration to AIX V4.1.4:

```
/etc/rc.ncs preserve
```

1.5 General TCP/IP Migration

Basic TCP/IP functionality includes the ability to use standard TCP/IP functions and communications. This includes connectivity using a default gateway/router, a hostname, an allocated IP address and subnet mask, and basic services such as ping.

This section covers the migration of a basic TCP/IP setup from AIX V3.2.5 to AIX V4.1.4. It also includes general considerations which apply to other more complex system setups.

1.5.1 General TCP/IP Migration Environment

A system was set up with AIX V3.2.5, a hostname, and a default gateway/router using the `smitty mktcpip` command.

This system was migrated to AIX V4.1.4 using the Migration Install method.

1.5.2 General TCP/IP Migration Planning

Due to some problems noted in the initial tests of Migration Install of AIX V4.1.4, we recommend that a copy of the /etc/rc.tcpip file is made before migration, say as /etc/rc.tcpip.325. This file, considered as a user file, will then be available after migration. The default gateway address should also be noted.

If you want to choose a Preservation or Overwrite installation, you should back up your TCP/IP configuration files separately. Which files you should back up will vary according to which services you are using and how the system is configured.

You should also be *very sure* that you have reliable backups of all of the data on your system, for both root VG and non-root VGs. One reason you may want to use Preservation, or Overwrite installation is to quickly build a clean AIX V4.1.4 system.

1.5.3 General TCP/IP Differences and Migration Experiences

During some tests of the migration process, some lines in /etc/rc.tcpip were reset back to the default AIX V4.1.4 settings. The only consequence noted was that some daemons which were set to automatically start up on reboot with AIX V3.2.5 were commented out; so they did not start up on reboot with AIX V4.1.4. This was only noted in some tests, and we were not able to determine the exact cause.

Attention!

After migration, be aware that adapter names may have changed.

During the migration to AIX V4.1.4, the device configuration is erased, then rebuilt. After this process, it is possible that network adapter names may have changed. For example, if before migration, the system had:

```
# lsdev -Ccadapter

tok0          Available 00-07        Token-Ring High-Performance Adapter (8fc8)
tok1          Available 00-05        Token-Ring High-Performance Adapter (8fc8)
```

Then, during the migration process which rebuilds the device database, the token-ring adapter in location 00-05 (adapter slot 5) would be discovered first, and would be named tok0. It would then inherit all of the device attributes that previously belonged to the adapter in location 00-07 (adapter slot 7).

After the migration, the system has:

```
# lsdev -Ccadapter

tok0          Available 00-05        Token-Ring High-Performance Adapter (8fc8)
tok1          Available 00-07        Token-Ring High-Performance Adapter (8fc8)
```

There are two ways to fix this problem. Use one of the following to fix any problems you have with swapped adapter names:

- Swap the cables so that the cable which was connected to the tok0 adapter on the system before migration (in location 00-07), is once again connected to the adapter known as tok0 after migration (now the adapter in location 00-05). This means that the adapter settings do not need to be changed.
- Reconfigure the adapter settings so that the settings for adapter tok0 after migration are the same as the settings which adapter tok1 had before the migration. This means the cabling does not have to be changed. Information about the adapter settings is stored in /tmp/reconfig1, /tmp/reconfig2, and /tmp/reconfig3 after the Migration Install of AIX V4.1.4. Use the information in these files, erase the existing adapter configurations, and then reconfigure the adapters to their original settings.

The migration process to AIX V4.1.3 has a problem with volume groups which have the `auto_on` attribute set to `y`. The migration process to AIX V4.1.4 will always set the `auto_on` attribute of non-root Volume Groups (VGs) to `y`. For more information, see the chapter about migrating AIX in *A Holistic Approach to AIX V4.1 Migration, Volume 1*.

Note that inetd information is no longer kept in the ODM. This is to eliminate the redundant configuration which was also kept in AIX files. The inetd daemon now relies only on the information in /etc/inetd.conf and /etc/services.

1.5.4 TCP/IP Interoperability - AIX V3.2.5 and AIX V4.1.4

There were no significant incompatibilities found in our tests between TCP/IP services in AIX V3.2.5 and AIX V4.1.4. Some minor differences were noted, and these are discussed in the relevant section.

A compatibility fileset, bos.compat.net, is provided with AIX V4.1.4. This fileset provides symbolic links for files which have been moved from their location in AIX V3.2. The compatibility fileset also provides some programs which are no longer a standard part of TCP/IP in AIX V4.1.4. The compatibility package is there to provide a migration environment for procedures and programs that rely on files and programs that are not a standard part of TCP/IP. In the future, these files, programs, and symbolic links may no longer be provided. If you find that you need to install the compatibility fileset, you should change your procedures, or notify the owners of any software that does not work without this fileset installed.

1.6 Migrating DNS

Human beings just aren't very good at remembering lots of numbers. That's why we have directories with file names, or folders with document names, instead of using "inode numbers" when we want to edit a document. The idea of subdirectories (or folders and documents) helps us to remember what we want to work on by giving that file a meaningful name.

It also gives us a way of organizing the names so that we do not see all the names all at once, so that we can navigate our way to a name by following a path or a branch down to our file.

A directory tree on a UNIX system or a PC is an example of a name space. A name space is used for more than files and directories. It is used for client/server computing and management by the Distributed Computing Environment (DCE) and also by DNS, the Domain Name System.

DNS is the way that host names are organized on the Internet using TCP/IP. Host names are used to look up or *resolve* the name we know a system as, and convert it to a TCP/IP address. All of the movement of data on a TCP/IP network is done using addresses, not host names so DNS is used to make it easy for humans to manage and work with the computer network.

If you have a site with many systems, you can use DNS to delegate the responsibility for naming systems to other people or sites. You can also reduce your administration workload by only having to update one server in case you want to change the address of a system.

1.6.1.1 DNS Domains

DNS uses a name space in a similar way to the directories and subdirectories we are used to. Instead of a ″/″ between names to show that we are going to the next level down, DNS uses a period or full stop like the one at the end of the last sentence.

The different DNS ″directories″ are called domains. Note that Network Information Service (NIS) also uses the term domain, but this is not related to a DNS domain. In this section, domain refers to a DNS domain, not to an NIS domain.

In the same way as / is the root directory for UNIX, DNS has . as the root of the name space. Unlike UNIX, if you leave out the full stop or period at the end of the DNS name, DNS will try various full or partial domain names for you. One other difference is that reading left to right, DNS goes from the lowest level to the highest, whereas the UNIX directory tree goes from the highest to the lowest.

For example, the domain ibm.com is a subdomain of the com domain. The domain austin.ibm.com is a subdomain of the ibm.com domain, and the .com domain.

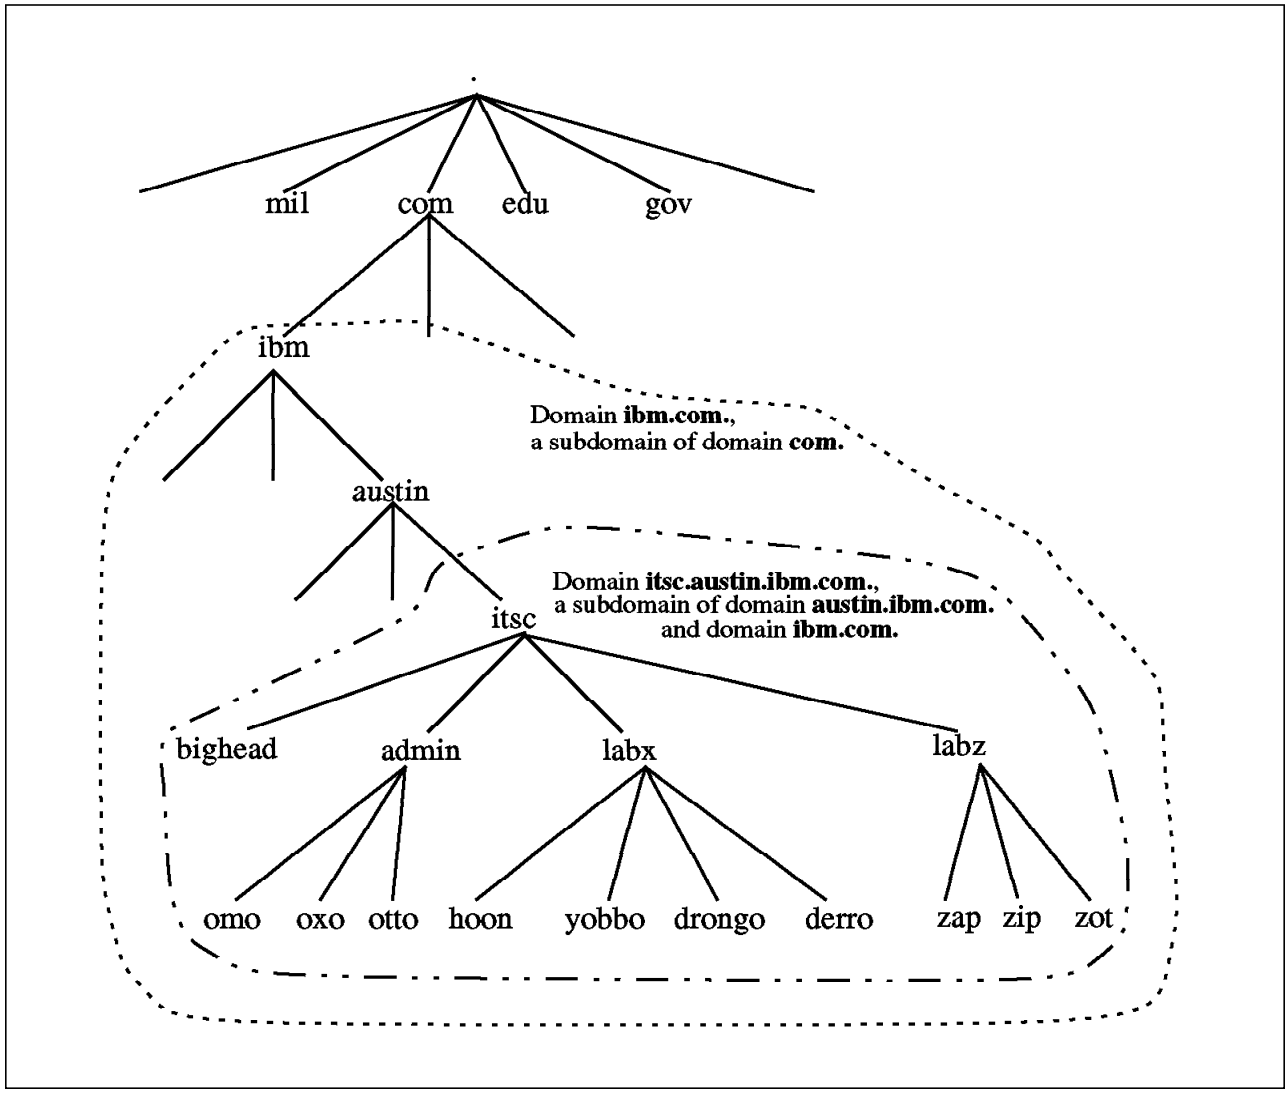

Figure 1. DNS Domains and Subdomains

You can set up a network without DNS. This uses a file called /etc/hosts on each system to define the mapping from names to TCP/IP addresses. Because each system has to have a copy of the /etc/hosts file, this becomes difficult to maintain for even a small number of systems. Even though setting up DNS is more difficult initially, the administrative workload for three or four workstations may be easier than with /etc/hosts. Maintaining a network of 20 or 30 workstations becomes just as easy as for three or four workstations.

When you set up DNS, you do not have to match your physical network to your DNS setup, but there are some good reasons why you should. Ideally, the primary and secondary name servers should be the systems which have the best connections to other domains and zones.

1.6.1.2 DNS Zones

DNS has the concept of domains and zones. A domain is a whole branch in the domain namespace. For example, austin.ibm.com is a domain. Any names of systems that end with .austin.ibm.com are within the austin.ibm.com domain. They are also in the ibm.com domain, and within the .com domain.

The term domain applies to the names we give to systems. The term zone applies to the servers responsible for resolving DNS names to network addresses. A zone is a part of a domain and is administered by a DNS name server. That name server is said to be the authority for that zone.

A zone could be a whole domain. We could have a name server that is the authority for the whole austin.ibm.com domain. Any requests to find the IP address of a host name ending in .austin.ibm.com will be handled by this name server. In this case, the server′s zone is the whole austin.ibm.com domain.

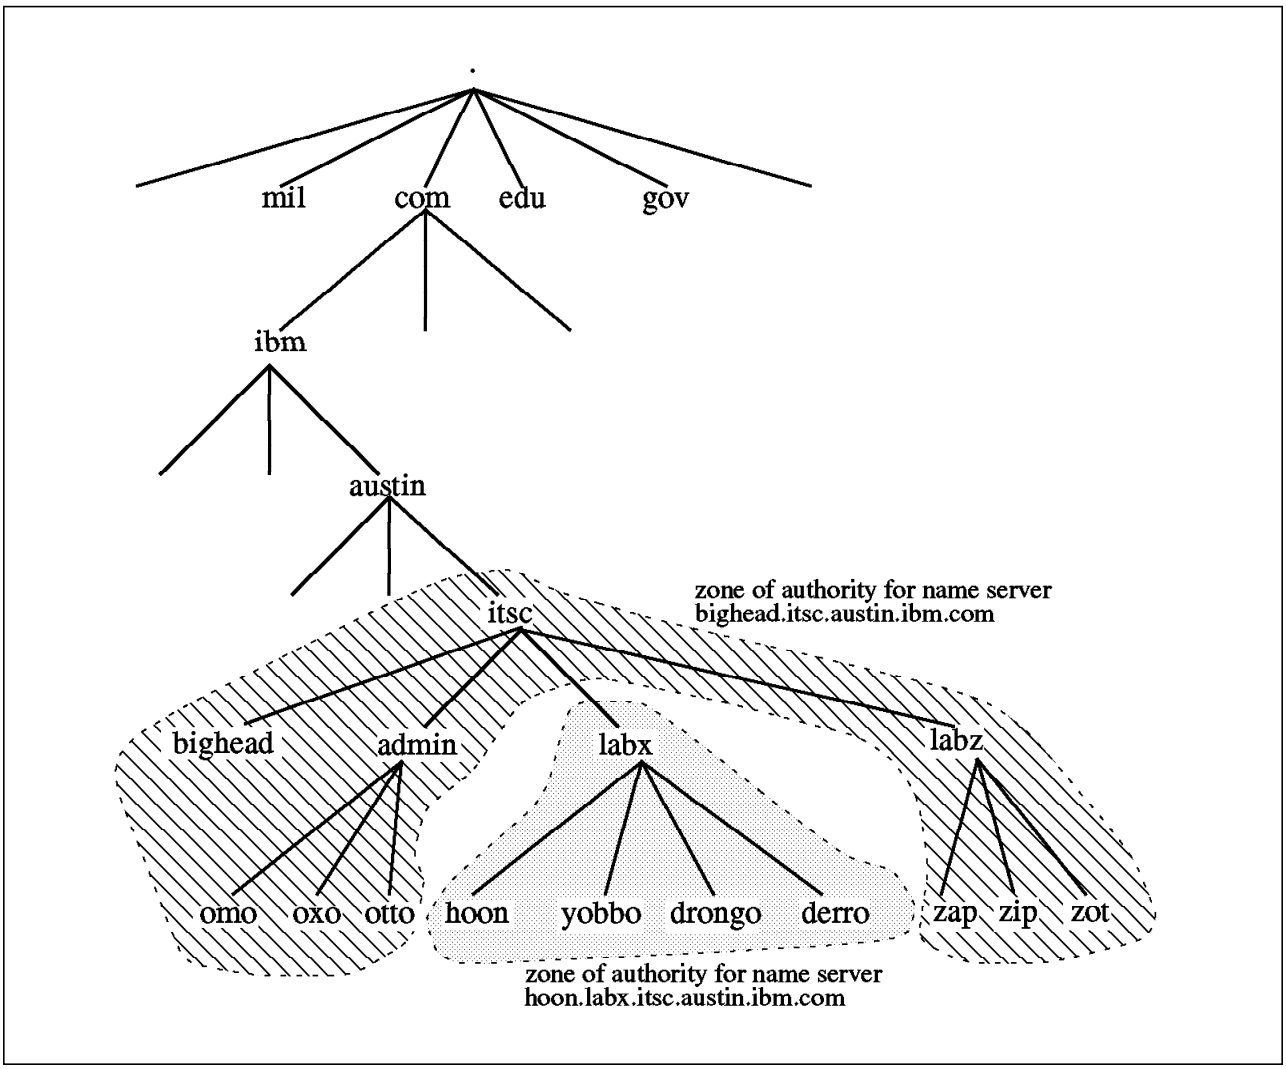

Figure 2. DNS Zones

A zone could also be *part* of a domain. Within the itsc.austin.ibm.com domain in Figure 2, we have one name server, called bighead.itsc.austin.ibm.com, which manages all host names ending with .itsc.austin.ibm.com, *except* those names which end with .labx.itsc.austin.ibm.com. The names ending with .labx.itsc.austin.ibm.com are managed by the name server hoon.itsc.austin.ibm.com. Each of these name servers manages host names within the itsc.austin.ibm.com domain, but each has a different zone which it controls. A name server usually has a host name within the zone it controls.

1.6.1.3 Primary and Secondary Name Servers

If a name server is responsible for the validity of names within a zone, it is called a primary master name server, or just a primary name server. There can only be one primary name server for any zone, but we can have secondary name servers which read and cache information from the primary name server. These secondary name servers can help reduce the workload on a primary name server and also allow name resolution if the primary name server is unavailable.

When the named daemon is started, it reads the files /etc/named.boot for DNS configuration information. The filename can be changed by invoking the named

daemon with the -b flag. If the named daemon is started by startsrc manually, the startup command to use the file /etc/named.my_left_boot would be:

```
# startsrc -a "-b /etc/named.my_left_boot" -s named
```

The entries in the DNS daemon's boot file (usually /etc/named.boot) tell the named daemon:

- which domains it will be a primary server for
- which domains it is a secondary server for
- which other name servers to contact to resolve names it does not know about
- which directory contains the configuration files
- which filenames are used as configuration files

In order to provide for redundancy in case a primary master name server is unavailable, DNS supports "secondary master" name servers. A primary master is responsible for information about name resolution within a zone. A secondary master uses information from the primary master and is available in case the primary master is not available.

Secondary name servers can also load their data from other secondary name servers. The system that the data is loaded from is known as a master server. So, a secondary name server loads its data from either a primary master server or a secondary master server. The process of loading data from a master server is known as a zone transfer.

The secondary name server can save a copy of the primary name server's name data on its own disk, or it can just keep a copy of the database in memory. Normally, secondary servers would be set up to save a copy of the name data on disk since this allows the secondary name server to be operational even if the primary name server is unavailable. The data kept on the secondary name server is kept synchronized by using the Serial value in the Start of Authority (SOA) record in the main data file. The secondary name server checks this regularly while it is running so that any changes are reflected in its own information. The time between checks is set by the Refresh value in the data file and is usually set to more than one hour. A secondary server which already finds its data saved on local disk will not reload the DNS information at startup. It waits until its timer has expired, checks the Serial on the primary name server and then reloads the data if necessary.

A secondary name server can also be a primary name server. In fact, it is recommended that a secondary server be a primary server for its own loopback address (127.0.0.1).

Attention!

There were problems running a secondary name server on a system migrated to AIX V4.1.4. See 1.6.4, "DNS Differences and Migration Experiences" on page 24 for more information.

If a system needs to resolve a name which is outside the zone of its name server, the request is still passed to the name server. It passes the request on to another name server (a forwarder) or to a root name server, which then hands that request down to a name server that has information about that zone. Forwarders insure that the ″top level″ name servers are not burdened with traffic that can be resolved at a lower level. The list of root name servers which will be contacted is held in the primary name server′s cache file (defined in the named.boot configuration file). This ″cache″ file is really not a cache, but it is called this for historic reasons. The list of forwarders to contact is defined in /etc/named.boot.

Figure 3 shows the relationship between various DNS configuration files on a system that is being used as a DNS primary or secondary name server.

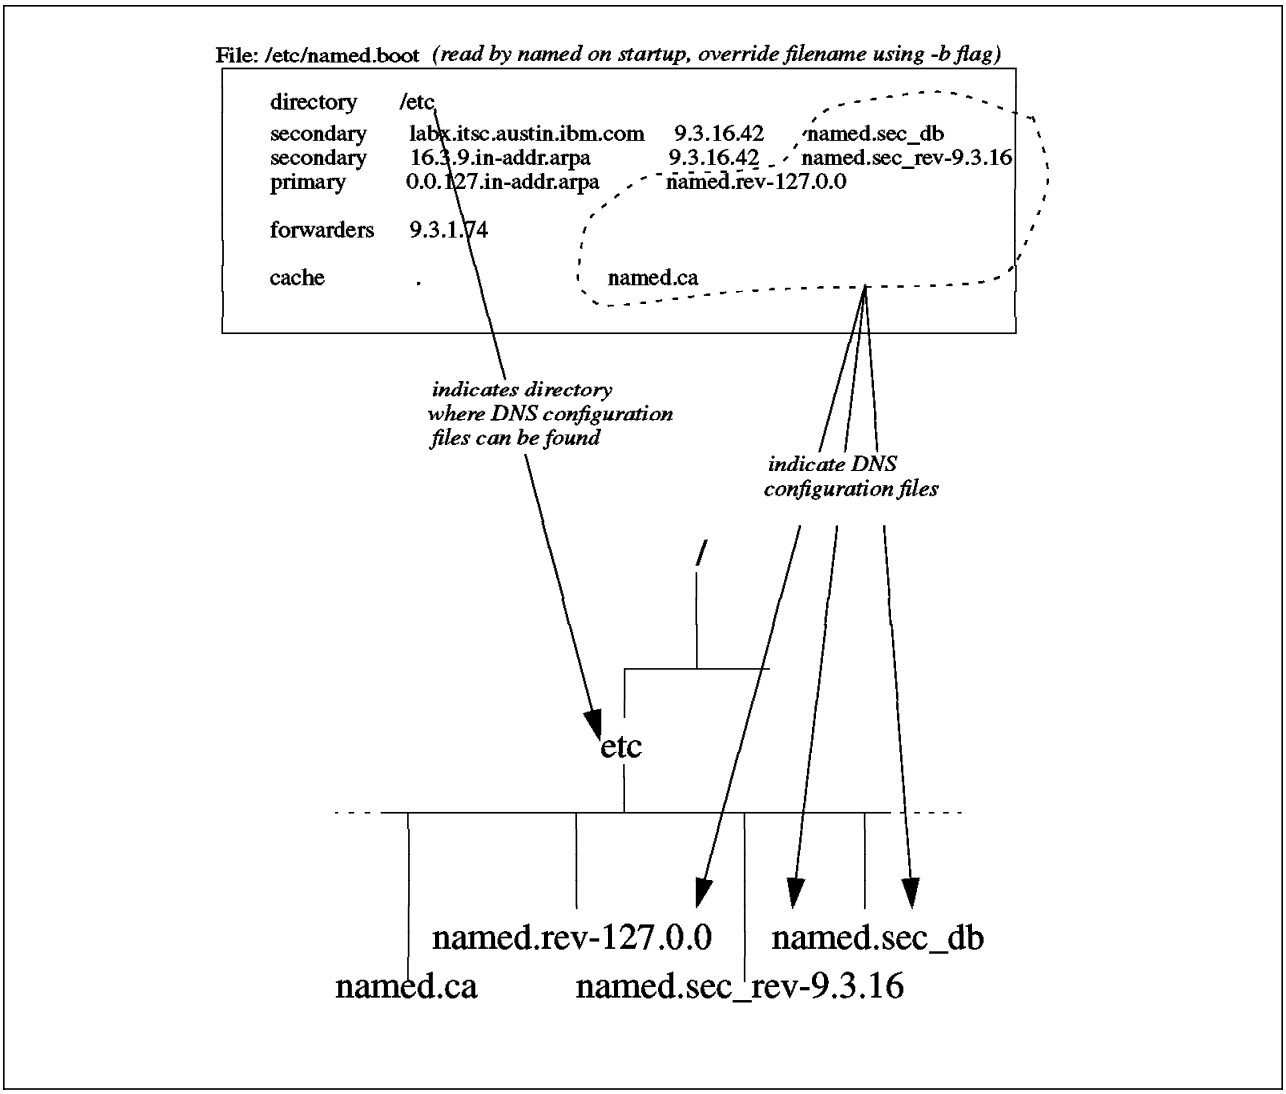

Figure 3. Files Used by DNS Name Server Daemon (named)

1.6.1.4 Name Resolvers

Name resolvers are services provided by the operating system that are used to translate system host names to IP addresses. Even on a name server, the name resolver services are used to pass the name resolution request onto the local name server daemon. The name resolvers are configured using the /etc/resolv.conf file although there are environment variables such as LOCALDOMAIN and HOSTALIASES which can override or modify the settings in this file. LOCALDOMAIN can override the default domain setting, and HOSTALIASES specifies a file to check instead of /etc/hosts.

If you modify /etc/resolv.conf, it will automatically be reread by the name resolver routines. There is no need to refresh or restart the named daemon.

On a client system, the /etc/resolv.conf file is set up so that the name resolver routines know which name server to contact and which domain names to search for a hostname. In /etc/resolv.conf, the domain directive is set to the client's domain name, and the name server directive is set to the network address of the name server to be used.

The search directive can be used instead of the domain directive. This allows a client to search for hostnames in two or more domains. If you often need to contact hosts in a particular domain that is not your own, this saves typing the whole domain name as well as the host name.

In the example below, a lookup of host ramjet fails because it is not within the labx.itsc.austin.ibm.com domain specified in /etc/resolv.conf. The host ramjet is actually in domain austin.ibm.com; so doing a lookup of ramjet.austin.ibm.com succeeds. The lookup of yobbo succeeds because yobbo is in the domain specified with the domain directive in /etc/resolv.conf.

```
# cat /etc/resolv.conf
domain labx.itsc.austin.ibm.com
nameserver 9.3.16.42

# ping ramjet
0821-062 ping: host name ramjet NOT FOUND

# ping ramjet.austin.ibm.com
PING ramjet.austin.ibm.com: (129.35.223.121): 56 data bytes
64 bytes from 129.35.223.121: icmp_seq=0 ttl=255 time=5 ms

----ramjet.austin.ibm.com PING Statistics----
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max = 5/5/5 ms

# ping yobbo
PING yobbo.labx.itsc.austin.ibm.com: (9.3.16.43): 56 data bytes
64 bytes from 9.3.16.43: icmp_seq=0 ttl=255 time=1 ms

----yobbo.labx.itsc.austin.ibm.com PING Statistics----
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max = 1/1/1 ms
```

Using the search directive instead of the domain directive in /etc/resolv.conf, we can add another domain to search. When we try to resolve a name, the name resolvers try to find it in both domains. In this example, we can find hostnames

in both the labx.itsc.austin.ibm.com domain and also in the austin.ibm.com domain.

```
# cat /etc/resolv.conf
search labx.itsc.austin.ibm.com austin.ibm.com
nameserver 9.3.16.42

# ping ramjet
PING ramjet.austin.ibm.com: (129.35.223.121): 56 data bytes
64 bytes from 129.35.223.121: icmp_seq=0 ttl=255 time=5 ms

----ramjet.austin.ibm.com PING Statistics----
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max = 5/5/5 ms

# ping yobbo
PING yobbo.labx.itsc.austin.ibm.com: (9.3.16.43): 56 data bytes
64 bytes from 9.3.16.43: icmp_seq=0 ttl=255 time=1 ms

----yobbo.labx.itsc.austin.ibm.com PING Statistics----
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max = 1/1/1 ms
```

The /etc/resolv.conf *must* exist on any AIX system that wants to use a DNS name server. This includes any system that is a primary or secondary name server. On a primary or secondary name server, the file must exist, but the contents are optional; it can be an empty file. A name server can be set up so that the resolvers follow the directives in the /etc/resolv.conf file. If /etc/resolv.conf is empty, the domain defaults to the domain of the local nameserver, and the nameserver defaults to the local system.

The default order in which various name resolvers are tried in AIX V4.1.4 is:

1. DNS, if it is considered to be enabled. If /etc/resolv.conf exists DNS is considered to be enabled.
2. NIS, if the ypbind daemon is running.
3. /etc/hosts on the local system if the NIS ypbind daemon is not running.

The default order of name resolution can be overridden using either the file /etc/netsvc.conf, or the environment variable NSORDER. The values which can be specified are:

- bind to specify DNS name resolution
- nis to specify NIS name resolution
- local to specify name resolution using /etc/hosts on the local system

User-defined name server routines can also be defined. For more information, see *Network Address Translation* in InfoExplorer.

If the string *=auth* is added directly following the specified name resolver service, then this facility is treated as the authoritative service. This means that if this service is running and the name resolution fails using this service, then no other services will be tried. However, if the authoritative service is not running, the next service will be tried.

The following example tries DNS/bind first, then NIS, and then the local /etc/hosts for name resolution. If DNS is running and the lookup fails, no other services are tried. If DNS is not running, then NIS is tried, followed by /etc/hosts.

```
# cat /etc/netsvc.conf

hosts = bind=auth,nis,local
```

The same effect could be achieved by setting the NSORDER environment variable (shown below for ksh).

```
# export NSORDER="bind=auth,nis,local"
```

1.6.2 DNS Migration Environment

The systems which were migrated were a DNS primary name server and a DNS secondary name server running on AIX V3.2.5. Both systems were migrated to AIX V4.1.4 by using the Migration Install method.

The DNS primary name server was migrated first to AIX V4.1.4, and then the AIX V3.2.5 secondary name server was tested for interoperability with the primary by adding a new name to the primary and forcing the secondary to reload its data files. The AIX V3.2.5 secondary name server was then migrated to AIX V4.1.4 and tested for proper operation.

A secondary name server running AIX V4.1.4 was tested with a primary name server running AIX V3.2.5 and also with a primary name server running AIX V4.1.4. In both cases, there was a problem reloading the secondary zone data.

For more information about interoperability, see 1.6.5, “DNS Interoperability - AIX V3.2.5 and AIX V4.1.4” on page 26.

Migrating a DNS secondary name server was similar to migrating a primary name server since the configuration files are also determined by the contents of the named.boot file.

1.6.3 DNS Migration Planning

Attention!

If you are planning to migrate a secondary name server to AIX V4.1.4, you should be aware that at the time of writing, there is a known problem with the secondary name server being unable to reload its zone information from a master server. Please check with your local IBM software support regarding this problem.

If you are planning to do a Migration Install of AIX V4.1.4, all name server data files in rootvg are preserved (unless they are in /tmp!). Name server data files in non-root Volume Groups (VGs) are preserved, but you should make sure that the VG is varied on and that the file system is mounted automatically after a reboot. There is more information about this in the *AIX* migration chapter in *A Holistic Approach to AIX V4.1 Migration, Volume 1*, SG24-4652.

In the case of any problems with the migration, you can recover your /etc/named.boot file from your backup tapes, and use the information in it to recover your DNS data files from your backup tapes. You will probably also want to recover /etc/resolv.conf and /etc/hosts.

If you choose to do a Preservation Install, then any DNS data files that are in the / (root), /usr, /var, or /tmp file systems will be lost. Only /home and any user-created file systems are preserved in the root VG. Note that the /etc directory is usually in the / (root) file system; so any DNS configuration files in /etc will be lost, too!

A New or Complete Overwrite Install (Overwrite Install) will destroy all files and file systems in rootvg, and in certain circumstances, may also add non-root VG disks to the rootvg. This could destroy additional file systems. You should make sure that your backup of both rootvg and non-rootvg file systems and data is complete and intact before trying any sort of migration. If you want to back up your DNS data files separately, back up /etc/named.boot (or your own named boot file), /etc/resolv.conf, /etc/hosts, and your DNS data files before starting a Preservation or Overwrite installation of AIX V4.1.

If your DNS data files are already in a file system on a non-root VG, then they will probably be available even after a Preservation, New or Complete Overwrite Install. Make sure the non-root VG is varied on, and the file systems are mounted automatically after migration completes. Since you have made a complete system backup before starting the migration process, all the DNS configuration files will be on your backup tapes if you need them.

If you wish to back up your DNS configuration files separately, you may wish to use the procedure below as a guide. Note that the filenames you should back up will vary depending on how your name server has been configured.

```
# cd /
# ls ./etc/named.boot ./etc/resolv.conf ./etc/hosts ./etc/named.* | \
> grep -v named.pid | backup -ivqf/dev/fd0

Backing up to /dev/fd0.
Cluster 9216 bytes (18 blocks).
Volume 1 on /dev/fd0
a ./etc/hosts
a ./etc/named.boot
a ./etc/named.boot
a ./etc/named.ca
a ./etc/named.db
a ./etc/named.rev-127.0.0
a ./etc/named.rev-9.3.16
a ./etc/resolv.conf
Backup finished on Thu Jan 25 16:44:57 CST 1996;
there are 18 blocks on 1 volumes.
```

You might notice that /etc/named.boot was backed up twice because it was specifically named and also because it matched the /etc/named.* pattern of names we've used for our DNS data files. Now we check to make sure the backup worked.

```
# restore -Tvqf/dev/fd0
```

In the example above, we did not save the /etc/named.pid file since this contains the process ID of a running named daemon and will be automatically created when the named daemon is started. When the migration is finished, the files can be restored by using:

```
# restore -xvqf/dev/fd0
```

1.6.4 DNS Differences and Migration Experiences

The problems and differences we found when migrating DNS from AIX V3.2.5 to AIX V4.1.4 were as follows:

- A secondary name server running AIX V4.1.4 could not reload its zone data from the primary server (either AIX V3.2.5 or AIX V4.1.4).
- After migration, the named daemon may not be started automatically.
- The AIX V4.1.4 named daemon does not allow a line continuation character (left bracket) until after the DNS administrator's E-mail address in the SOA record. This was allowed in AIX V3.2.5.
- The AIX V4.1.4 name resolvers now support more than two domains in the search directive in /etc/resolv.conf.

A secondary name server running AIX V4.1.4 could not reload its zone information from the DNS primary name server. This is due to the fact that a link is missing from the file /etc/named-xfer to /usr/sbin/named-xfer. To fix this problem, use the following command:

```
ln -s /usr/sbin/named-xfer /etc/named-xfer
```

If you wish to receive any AIX updates related to this problem, quote APAR IX55267 (PTF U441125) to your IBM software support contact.

After any installation of AIX V4.1.4, you may need to change the name server daemon to start automatically. In some cases, we noted that a Migration Install resulted in a /etc/rc.tcpip with a disabled named. Enabling named can be done with the smit stnamed fastpath and by then selecting the **BOTH** option.

```
                        Start Using the named Subsystem

Move cursor to desired item and press Enter.

  NOW
  Next System RESTART
  BOTH



F1=Help             F2=Refresh          F3=Cancel           F8=Image
F9=Shell            F10=Exit            Enter=Do
```

As an alternative, use a text editor to edit the /etc/rc.tcpip file, and remove the # from the beginning of the line that reads

```
# start /usr/sbin/named "$src_running"
```

The named daemon in AIX V4.1.4 is more fussy about the format of the DNS data files than the named daemon in AIX V3.2.5. After migration, the named would not start since it could not process the data files with SOA records in them. To find out what the problem was, we looked at the debug output file /var/tmp/named.run, after starting the named daemon with the following command:

```
# startsrc -s named -a "-d 11"
```

The reason for the failure of the named daemon was that the original AIX V3.2.5 data files had SOA records that used a left bracket "(" after the SOA system name and before the mail address of the administrative contact to show that this line was continued until a matching right bracket ")" ended the line. The AIX V3.2.5 data file reads:

```
@  IN  SOA  hoon.labx.itsc.austin.ibm.com. (
                    a948r5.itsorus.itsc.austin.ibm.com ; DNS administrator
                    199601261436 ; Serial
                    3600         ; Refresh after 1 hour
                    300          ; Retry after 300 seconds
                    360000       ; Expire after 100 hours
                    86400  )     ; Minimum TTL of 1 day
```

The AIX V4.1.4 named daemon required that the left bracket indicating a line continuation be placed *after* the DNS administrator's E-mail address. This is shown in the updated file below that is used in AIX V4.1.4.

```
@  IN  SOA  hoon.labx.itsc.austin.ibm.com. a948r5.itsorus.itsc.austin.ibm.com (
                    199601261436 ; Serial
                    3600         ; Refresh after 1 hour
                    300          ; Retry after 300 seconds
                    360000       ; Expire after 100 hours
                    86400  )     ; Minimum TTL of 1 day
```

The AIX V4.1.4 named daemon supports multiple searches for hostnames in different domains by using the search directive in the /etc/resolv.conf file. In AIX V3.2.5, you could only have two domains specified after the search directive in your /etc/resolv.conf file, and you also had to manually edit the /etc/resolv.conf file if you wanted to use the search directive. With AIX V4.1.4, the search directive is now supported by SMIT (fastpath smit domainsearch).

In AIX V3.2.5, if you were in domain labx.itsc.austin.ibm.com, the name resolvers would resolve drongo.labx as well as drongo or drongo.labx.itsc.austin.ibm.com. If you want this same behavior in AIX V4.1.4, then you should use the *search* parameter in /etc/resolv.conf.

1.6.5 DNS Interoperability - AIX V3.2.5 and AIX V4.1.4

A DNS client running on a system with AIX V3.2.5 was tested with a primary name server running on another system with AIX V4.1.4. Name resolution was found to work properly.

A secondary DNS name server running AIX V3.2.5 was tested with a primary name server system running AIX V4.1.4. The secondary name server reloaded its secondary zone correctly from the primary name server and worked correctly.

Attention!

There were problems running a secondary name server on AIX V4.1.4. For more information, see 1.6.4, “DNS Differences and Migration Experiences” on page 24.

1.7 Migrating r Commands

The `r` commands are a convenient and powerful set of commands which can be used between TCP/IP networked systems. The commands include:

`rlogin` Allows users to log in to another system

`rsh` Executes a command on another system

`rexec` Executes a command on another system and can automatically log in

`rcp` Copies file between different systems using a syntax similar to the UNIX `cp` command.

All of the `r` commands use the /etc/hosts.equiv and $HOME/.rhosts files to determine if the user attempting to use a service is automatically authenticated or if the user must enter a username and/or a password.

If non-root users attempt to use an `r` command, AIX first checks the /etc/hosts.equiv file, and then the file $HOME/.rhosts. Here, $HOME stands for the home directory of each user. For example, if a user bertha attempts to log in, the /etc/hosts.equiv file is checked first, and then the .rhosts file in the home directory of the local user bertha (usually /home/bertha/.rhosts) is checked. For root users attempting to use one of the `r` commands, only the /.rhosts file is checked.

The files contain specific hosts, specific users, or specific users on specific hosts, which are allowed (or not allowed) to use the `r` commands. The routines which perform the authorization are also NIS aware, and you can use NIS netgroups and NIS maps for both hosts and users in these files. The files are processed line-by-line until a match is found which either allows or disallows access.

The format of the $HOME/.rhosts file and /etc/hosts.equiv file is the same.

The following file disallows user paulk on system yobbo.labx.itsc.austin.ibm.com, allows root on yobbo.labx.itsc.austin.ibm.com, allows all users on system hoon.labx.itsc.austin.ibm.com, and disallows all users on drongo.labx.itsc.austin.ibm.com.

```
yobbo.labx.itsc.austin.ibm.com -paulk
yobbo.labx.itsc.austin.ibm.com root
hoon.labx.itsc.austin.ibm.com +
-drongo.labx.itsc.austin.ibm.com
```

You can see that since a file in a user's home directory can provide authorization for access to your system, it is a very vulnerable point of security. It is also difficult to manage. As a result, many system administrators disable the r commands for the whole system by commenting out the lines starting with shell, login and exec in the /etc/inetd.conf file. These entries correspond to the daemons rshd, rlogind, and rexecd which perform the r commands on the local system.

The services which the various daemons provide are:

- rlogind provides services for rlogin.
- rshd provides services for rsh and rcp.
- rexecd provides services for rexec.

1.7.1 r Command Migration Environment

The r commands were tested on systems running AIX V3.2.5 and AIX V4.1.4. These systems were also running in a DNS (Domain Name System) and NFS (Network File System) environment.

Commands tested were:

- rlogin
- rsh
- rexec

The server daemons for the r commands are started by default in AIX V3.2.5 and AIX V4.1.4 (see the inetd "super-daemon" configuration file /etc/inetd.conf).

1.7.2 r Command Migration Planning

There was no requirement to save and restore files since the /.rhosts and /etc/hosts.equiv files were not affected by the Migration Install.

1.7.3 r Command Differences and Migration Experiences

There were no differences noted between the r commands on AIX V3.2.5 and AIX V4.1.4.

1.7.4 r Command Interoperability - AIX V3.2.5 and AIX V4.1.4

There were no problems experienced in interoperability between the r commands tested on the two different versions of AIX.

1.8 Migrating NFS

Attention!

A problem with AIX V4.1.4 migration means that some NFS and NIS files must be manually recovered from the AIX distribution tape or CD and manually merged with the customized AIX V3.2.5 files left on the system. See 1.8.3, “NFS Differences and Migration Experiences” on page 31 for more information.

NFS provides the ability to read and write to files which are located physically on another computer system.

NFS allows a directory which is in a file system on another system to be accessed on the local system. To a user, the remote directory and the files in the remote directory appear to be local files. The process of making a remote directory accessible locally involves *mounting* the directory in a similar way to mounting a local file system.

With NFS, any system can be both a client and a server. A server is a system that is set up to provide access to its local directories and files. A client is a system that is accessing the directories and files on another system. One system can access the files and directories on another system and, at the same time, *export* its own directories and files to make them available to other systems.

These functions are provided by a combination of the AIX kernel and NFS daemon processes. Multiple NFS daemons are used to service multiple requests at once.

When a user on a system accesses a directory or a file which is NFS mounted, this request is sent to the server across the TCP/IP network. A `biod` daemon on the client improves the apparent performance of NFS by pre-fetching information into a cache and by doing write-behind for the client. Write-behind means that when a client writes, the information is written to a buffer, and the client is told the write succeeded before the information is actually sent across the network to the NFS server.

The NFS client system passes the request across the TCP/IP network to an `nfsd` NFS server daemon on the server system. This daemon requests the actual data from the physical file system and sends it back to the client that requested it.

The process of mounting a remote file system onto a system can be done automatically or interactively by using the `mount` command. If you want to mount an NFS directory, information about the NFS mount can be placed into the /etc/filesystems file. The /etc/filesystems file contains information about file systems which can be mounted on this system. The /etc/filesystems file is equivalent to the /etc/fstab file found on systems based on BSD UNIX. The first stanza in the extract of /etc/filesystems seen below is for a local file system. The second stanza is for the NFS file system, /usr/local/hoon_files, which will be mounted from a system called hoon.labx.itsc.austin.ibm.com.

```
/usr/local/yobbo_files:
        dev             = /dev/lv00
        vfs             = jfs
        log             = /dev/hd8
        mount           = true
        check           = true
        options         = rw
        account         = false

/usr/local/hoon_files:
        dev             = "/usr/local/hoon_files"
        vfs             = nfs
        nodename        = hoon.labx.itsc.austin.ibm.com
        mount           = true
        options         = bg,soft,intr,retry=3
        account         = false
```

If the mount parameter is set to true, the system will attempt to mount the NFS directory when the system boots. This can delay the boot process since the system retries to mount the file system even if the NFS server is not available. You should be careful in understanding the implications of setting the mount parameter to true.

A client system is not allowed to mount any file system or directory available on a system which is configured to be an NFS server. On the server system, file systems or directories must be *exported* to make them available to clients.

The file /etc/exports contains a list of file systems and directories which are to be made available and a list of which systems are allowed to access these file systems and directories. The file systems and directories can be exported read/write or read-only to the client systems. In the example of /etc/exports shown below, /usr/local/yobbo_files is exported with read and write access allowed to everyone. Directory /usr/local/my_files is exported read-only, and the only host allowed access is hoon.labx.itsc.austin.ibm.com.

```
# cat /etc/exports

/usr/local/yobbo_files -
/usr/local/my_files -ro,access=hoon.labx.itsc.austin.ibm.com
```

NFS does not automatically allow root (superuser) level of access to exported files and directories. If this was done, security could be compromised. By default, the root user on a client system is given access to the server's files as the user nobody. The server can allow root access from specific clients, but this is discouraged for security reasons. Users which are unknown on the server system are also mapped to the user nobody.

NFS can be used in conjunction with Network Information System (NIS) in order to provide consistent access for users across multiple systems. NFS allows the same directories to be visible across many systems, and NIS allows user and system information to be kept consistent across many systems.

NFS can be configured by using SMIT with the fastpath `smit nfs`. It can also be configured by manually editing the files /etc/rc.nfs, /etc/filesystems and

/etc/exports. Using SMIT is recommended instead of manual editing. The file /etc/rc.nfs is the script used to start up NFS and NIS when the system boots, and the script /etc/nfs.clean is used to shut down NFS and NIS.

1.8.1 NFS Migration Environment

The environment used to test NFS migration consisted of two systems:

- *hoon*- A system running NFS, also acting as a DNS primary name server. This system had an NFS exported directory, and NFS mounted a directory from the second system (yobbo).
- *yobbo*- A system running NFS, also acting as a DNS secondary name server. This system had an NFS exported directory, and NFS mounted a directory from the first system (hoon).

First, both systems were tested together for correct operation at AIX V3.2.5. This involved checking that the NFS daemons were started, and that both systems could act as NFS clients and NFS servers to the other system.

Then, hoon was migrated to AIX V4.1.4 using the Migration Install method. Again, both systems were tested for correct operation.

Finally, yobbo was migrated to AIX V4.1.4 by using the Migration Install method. Both systems were tested to check that NFS read and write access was working.

The same tests were also performed in an environment which included NIS.

1.8.2 NFS Migration Planning

The general principles of planning to migrate your system to AIX V4.1.4 should be observed. For details on this, see *A Holistic Approach to AIX V4.1 Migration, Volume 1*, SG24-4652.

Based on the test experiences, there are no specific requirements for planning to migrate an NFS client or NFS server if you are going to use the Migration Install method.

Attention!

Due to problems in the Migration Install logic, you should follow the procedures in 1.8.3, “NFS Differences and Migration Experiences” on page 31 to successfully complete NFS migration.

NFS may appear to work correctly without doing this step, but this step is essential to make sure that your system will continue to work correctly with future releases of AIX.

If you are considering installing AIX V4.1.4 using the Preservation or Overwrite methods of installation, you will have to restart NFS on the system by using SMIT or by editing the configuration files. You should be prepared to recover the following files:

- /etc/filesystems which contains definitions for directories which are to be NFS mounted from other systems. This file should be restored without overwriting the default AIX V4.1.4 file, and then the NFS specific stanzas should be manually merged with the default AIX V4.1.4

- /etc/exports which defines which local directories are to be made available to other systems.
- /etc/rc.nfs which is the startup script for NFS and NIS. This file should be restored without overwriting the default AIX V4.1.4 file and then manually merged with the default AIX V4.1.4 file.
- /etc/nfs.clean which is the shutdown script for NFS and NIS. This file should be restored without overwriting the default AIX V4.1.4 file and then manually merged with the default AIX V4.1.4 file.

You should also note which NFS daemons are automatically started on your AIX V3.2.5 system by using the command:

```
lssrc -g nfs
```

1.8.3 NFS Differences and Migration Experiences

NFS in AIX V4.1.4 includes the ability for the NFS client to perform dynamic/adaptive retransmissions for its requests to the server.

After performing a Migration Install of AIX V4.1.4, a problem was noted with the *post_i* shell scripts which are invoked after product installation. The filesets affected are bos.net.nfs.client and bos.net.nis.client.

You must perform some manual steps after the AIX V4.1.4 Migration Install in order to complete the migration of NIS and NFS.

This problem has the following effects:

- AIX V3.2.5 customized NIS and NFS configuration files that should have been replaced by AIX V4.1.4 files are left in place.
- AIX V3.2.5 customized NIS and NFS configuration files which should be left in the /lpp/save.config directory after migration are erased.
- AIX V4.1.4 default NIS and NFS configuration files which should be left in the /lpp/save.config directory after migration are erased.

Some configuration files are meant to be replaced by the default AIX V4.1.4 configuration files during the Migration Install of AIX V4.1.4. These files are meant to be treated as *user_merge*, and instead are treated as *preserve* due to a problem in the post_i shell scripts for NIS and NFS. Copies of both the original (customized) AIX V3.2.5 files and the default AIX V4.1.4 files should be left in the /lpp/save.config directory.

The way that the AIX V4.1.4 migration process was *supposed* to work was that after migration, you would:

- Manually merge the customized AIX V3.2.5 file with the AIX V4.1.4 file by updating the relevant parts of the AIX V4.1.4 default file. Both files should have been available in /lpp/save.config.
- Replace the default AIX V4.1.4 file with the merged AIX V4.1.4 file.

Instead, the customized AIX V3.2.5 files are left in place after migration. Due to the high compatibility of AIX V4.1.4 with AIX V3.2.5, the AIX V3.2.5 files caused no problems with the TCP/IP subsystems we tested in AIX V4.1.4.

Attention!

Despite the lack of problems using the AIX V3.2.5 files, you should perform the manual merging process to avoid problems with future releases of AIX.

Even though AIX V4.1.4 NIS and NFS can successfully use the AIX V3.2.5 files, you should still do the manual merge. This is because future versions of AIX may rely on having the AIX V4.1 files. Future versions may also not include the same level of compatibility with AIX V3.2.5.

The only filesets and files affected by this problem were:

- bos.net.nfs.client
  - /etc/nfs.clean
  - /etc/rc.nfs
- bos.net.nis.client
  - /var/yp/Makefile
  - /var/yp/updaters

1.8.3.1 Recovering the Default AIX V4.1.4 Files

You will need to recover the default AIX V4.1.4 files from the AIX V4.1.4 distribution media (tape or CD) once the Migration Install has completed. The post_i shell script changed migration keywords for the affected files from user_merge to preserve. This means the default AIX V4.1.4 files are not available on the system after migration. For more information, see 1.4.1, “Files and Filesets” on page 9.

1.8.3.2 Recovering Default Files from CD

If you have a CD with AIX V4.1.4, follow this example to recover the default AIX V4.1.4 files to /tmp/org414/nfs and /tmp/org414/nis.

```
# mkdir /tmp/org414

# cd /tmp/org414

# mkdir /tmp/cd414

# crfs -v cdrfs -p ro -d'cd0' -m'/tmp/cd414' -A'no'

# mount /tmp/cd414

# restore -xvqf /tmp/cd414/usr/sys/inst.images/bos.net \
 ./usr/lpp/bos.net/inst_root/etc/rc.nfs \
 ./usr/lpp/bos.net/inst_root/etc/nfs.clean \
 ./usr/lpp/bos.net/inst_root/var/yp/Makefile \
 ./usr/lpp/bos.net/inst_root/var/yp/updaters

New volume on /tmp/cd414/usr/sys/inst.images/bos.net:
Cluster size is 51200 bytes (100 blocks).
The volume number is 1.
The backup date is: Tue Oct  3 12:16:41 CDT 1995
Files are backed up by name.
The user is BUILD.
x         2466 ./usr/lpp/bos.net/inst_root/etc/nfs.clean
x         4201 ./usr/lpp/bos.net/inst_root/etc/rc.nfs
x        10838 ./usr/lpp/bos.net/inst_root/var/yp/Makefile
x          469 ./usr/lpp/bos.net/inst_root/var/yp/updaters
The total size is 17974 bytes.
The number of restored files is 4.

# mkdir /tmp/org414/nfs

# cp /tmp/org414/usr/lpp/bos.net/inst_root/etc/* /tmp/org414/nfs

# mkdir /tmp/org414/nis

# cp /tmp/org414/usr/lpp/bos.net/inst_root/var/yp/* /tmp/org414/nis

# rmfs /tmp/cd414
```

1.8.3.3 Recovering Default Files from Tape

If you have an AIX V4.1.4 distribution tape (*not* a mksysb backup tape), use the following example as a guide for restoring the default AIX V4.1.4 files to /tmp/org414/nfs and /tmp/org414/nis. If you have an 8mm tape drive, use the following command before trying to restore the files to make sure that the block size of the tape device is set to 512 bytes/block:

```
lsattr -E -l rmt0
chdev -a block_size=512 -l rmt0
```

You will have to reset the tape device to its original block_size setting after you have finished restoring files (note the block_size setting from the lsattr command above). The following procedure to restore the AIX V4.1.4 files requires approximately 11 MB free in the /tmp file system.

```
# mkdir /tmp/org414

# cd /tmp/org414

# tctl -f /dev/rmt0.1 fsf 2

# dd if=/dev/rmt0 bs=100b | grep "bos.net {"
01:005:10539008 4 R I bos.net {
13+1 records in.
13+1 records out.

# tctl -f /dev/rmt0.1 fsf 4

# dd if=/dev/rmt0 of=/tmp/org414/bos.net.tape bs=1000b

# restore -xvqf /tmp/org414/bos.net.tape \
 ./usr/lpp/bos.net/inst_root/etc/rc.nfs \
 ./usr/lpp/bos.net/inst_root/etc/nfs.clean \
 ./usr/lpp/bos.net/inst_root/var/yp/Makefile \
 ./usr/lpp/bos.net/inst_root/var/yp/updaters

# mkdir /tmp/org414/nfs

# cp /tmp/org414/usr/lpp/bos.net/inst_root/etc/* /tmp/org414/nfs

# mkdir /tmp/org414/nis

# cp /tmp/org414/usr/lpp/bos.net/inst_root/var/yp/* /tmp/org414/nis
```

The number 005 in the string 01:005:10539008 is the sequence number of the tape file which contains bos.net. If you do not see 005, then use whatever number you see in the second position in the output of the command

```
# dd if=/dev/rmt0 bs=100b | grep "bos.net {"
```

and subtract 1 from it. Use this result instead of 4 in the `tctl -f /dev/rmt0.1 fsf 4` command. The `fsf 4` in the `tctl` command tells the tape drive to skip four End of File (EOF) tape markers and leaves the tape positioned at the fifth tape file.

1.8.3.4 Recovering Default Files from NIM Master

If you have a NIM server, use the following example as a guide to restore the default AIX V4.1.4 files to /tmp/org414/nfs and /tmp/org414/nis.

```
# lsnim -t lpp_source
labx_simages     resources       lpp_source

# lsnim -l labx_simages
labx_simages:
   class       = resources
   type        = lpp_source
   server      = master
   location    = /export/labx_simages         <note this directory>
   comments    = install images of course
   alloc_count = 0
   Rstate      = ready for use
   prev_state  = unavailable for use
   simages     = yes

             <use the directory you noted in the command below>

# ls /export/labx_simages/bos.net.*
/export/labx_simages/bos.net.usr.4.1.3.0

# mkdir /tmp/org414

# cd /tmp/org414

# restore -xvqf /export/labx_simages/bos.net.usr.4.1.3.0 \
 ./usr/lpp/bos.net/inst_root/etc/rc.nfs \
 ./usr/lpp/bos.net/inst_root/etc/nfs.clean \
 ./usr/lpp/bos.net/inst_root/var/yp/Makefile \
 ./usr/lpp/bos.net/inst_root/var/yp/updaters

New volume on /export/labx_simages/bos.net.usr.4.1.3.0:
Cluster size is 51200 bytes (100 blocks).
The volume number is 1.
The backup date is: Tue Oct  3 12:16:41 CDT 1995
Files are backed up by name.
The user is BUILD.
x         2466 ./usr/lpp/bos.net/inst_root/etc/nfs.clean
x         4201 ./usr/lpp/bos.net/inst_root/etc/rc.nfs
x        10838 ./usr/lpp/bos.net/inst_root/var/yp/Makefile
x          469 ./usr/lpp/bos.net/inst_root/var/yp/updaters
The total size is 17974 bytes.
The number of restored files is 4.

# mkdir /tmp/org414/nfs

# cp /tmp/org414/usr/lpp/bos.net/inst_root/etc/* /tmp/org414/nfs

# mkdir /tmp/org414/nis

# cp /tmp/org414/usr/lpp/bos.net/inst_root/var/yp/* /tmp/org414/nis
```

1.8.3.5 Merging Customized and Default Files

You now have the missing AIX V4.1.4 files in /tmp/org414/nfs and /tmp/org414/nis.

For each of the AIX V4.1.4 files you have just restored, you can make a backup of the AIX V4.1.4 default file and a backup of the customized AIX V3.2.5 version of the file currently being used on the running system:

```
# cp /tmp/org414/nfs/rc.nfs /tmp/org414/nfs/rc.nfs.test

# cp /etc/rc.nfs /tmp/org414/nfs/rc.nfs.325
```

You can now compare the customized AIX V3.2.5 files with the AIX V4.1.4 default files. You can print out the files and compare them, or you may want to use the `diff` command to assist you.

```
# diff /tmp/org414/nfs/rc.nfs.test /tmp/org414/nfs/rc.nfs.325

...
103,105c36,38
< #if [ -x /usr/bin/domainname ]; then
< #     /usr/bin/domainname ibm
< #fi
---
> if [ -x /usr/bin/domainname ]; then
>       /usr/bin/domainname labx
> fi
---
119,120c52,53
< #if [ -x /usr/lib/netsvc/yp/ypbind ]; then
< #     start ypbind /usr/lib/netsvc/yp/ypbind
---
> if [ -x /usr/etc/ypbind ]; then
>      startsrc -s ypbind
...
```

You now have to manually merge the information in the AIX V3.2.5 version of the file into the AIX V4.1.4 version of the file.

The aim of this step is to change the AIX V4.1.4 file by adding any customization information that may have been added to the AIX V3.2.5 version of the file. The end result should be an AIX V4.1.4 file that achieves the same result as the AIX V3.2.5 file. After merging, the AIX V4.1.4 file should be able to be used to start the same services and set up the same parameters as the AIX V3.2.5 file. Note that the locations and names of the services may have changed, as may the method of starting that service.

In the previous example, we can see that the NIS domain name needs to be changed to labx and the lines should be uncommented. Typically, the only thing you will need to do is uncomment lines applicable to the TCP/IP services you are using, and add customization information for these services.

You now have a choice of using `vi` to perform the merge of the two versions of the file, or you may also choose to perform the customization by using the appropriate SMIT menus. If you choose to use the SMIT menus to do the customization, you should copy the `.test` version of the file to its proper place (`cp /tmp/org414/nfs/rc.nfs.test /etc/rc.nfs`), and *then* do the customization using SMIT. In the example above, the NIS domain name could either be set by modifying the /tmp/org414/nfs/rc.nfs.test file using `vi`, or by using the SMIT fastpath `smit chypdom`.

Attention!

Do not change the paths and commands in the test (AIX V4.1.4) version of the file to match those in the running (AIX V3.2.5) files. Changing the paths and commands would reset the commands and paths back to AIX V3.2.5 standards and undo all the good work you have been doing to create a proper customized AIX V4.1.4 file!

As an example, in rc.nfs, you should leave the AIX V4.1.4 syntax of `start ypbind /usr/lib/netsvc/yp/ypbind` and associated lines as they are, and just remove the # from the line to uncomment it if you want to use this service.

```
# vi /tmp/org414/nfs/rc.nfs.test
...
                <make your changes, then exit vi>
```

Finally, you copy the modified version of the file to the correct directory:

```
# cp /tmp/org414/nfs/rc.nfs.test /etc/rc.nfs
```

Repeat this merge process for the four files which need to be updated to AIX V4.1.4 versions. If you are using NIS, you should do this merge process for all four affected files, that is, all of the files that apply to NFS *and* NIS. If you are only using NFS, you can just perform a merge on /etc/nfs.clean and /etc/rc.nfs.

To test the changes for:

- /etc/nfs.clean: If you are using NFS or NIS, use the command `/usr/bin/ksh /etc/nfs.clean`. This will stop NFS and NIS immediately.
- /etc/rc.nfs: If you are using NFS or NIS, shut the system down and reboot, then check for correct NFS and NIS operation.
- /var/yp/Makefile: If this system is an NIS master server, use the commands `cd /var/yp` and then `make`.
- /var/yp/updaters: Used in a system running secure NFS. Check that public keys can be updated and that secure NFS is working correctly.

1.8.4 NFS Interoperability - AIX V3.2.5 and AIX V4.1.4

NFS clients and servers running on either AIX V3.2.5 or AIX V4.1.4 were found to work properly with each other.

1.9 Migrating NIS

Attention!

A problem with AIX V4.1.4 migration means that some NFS and NIS files must be manually recovered from the AIX distribution tape or CD and manually merged with the customized AIX V3.2.5 files left on the system. See 1.8.3, “NFS Differences and Migration Experiences” on page 31 for further information.

NIS is a useful tool to assist in administering a large number of systems. The main purpose of NIS is to distribute up to date information from AIX files used for user management, system management, and network management. NIS can also be used to distribute information from your own files.

NIS is most commonly used to keep user names, user IDs, passwords, group names and group IDs consistent across many systems.

1.9.1.1 NIS Maps and Servers

NIS does not distribute the actual files containing the data. It uses the information in the files to build an NIS *map*, which is really a database file created and accessed by NIS clients. NIS uses the dbm database supplied as standard with AIX. Note that dbm is a *very* simple database and is not designed to provide the facilities, robustness and performance of commercial Relational Database Management System (RDBMS) products.

The information in the NIS maps is kept on a master server, which controls the information. Additional slave servers can hold copies of the information controlled by the master server; so performance and availability of information is improved. The availability of this information is crucial since it can include such things as hostname to IP mapping (/etc/hosts), user names and passwords (/etc/passwd, /etc/security/passwd). If the master server is not available and there is no slave server, a network of systems can be completely disrupted, with no systems operational.

The following table shows the NIS maps that are created from AIX files and other information on the master server that can be administered by NIS.

Table 2 (Page 1 of 2). NIS Default Maps.

Map	NIS Nickname	Files Used to Create Map
passwd.byname	passwd	/etc/passwd, /etc/security/passwd
passwd.byuid		
group.byname	group	/etc/group
group.bygid		
hosts.byaddr	hosts	/etc/hosts
hosts.byname		
ethers.byaddr	ethers	/etc/ethers
ethers.byname		
networks.byaddr	networks	/etc/networks
networks.byname		
rpc.bynumber		/etc/rpc

Table 2 (Page 2 of 2). NIS Default Maps.		
Map	**NIS Nickname**	**Files Used to Create Map**
services.byname	services	/etc/services
protocols.byname	protocols	/etc/protocols
protocols.bynumber		
netgroup		/etc/netgroup
netgroup.byhost		
netgroup.byuser		
bootparams		/etc/bootparams
mail.aliases	aliases	/etc/aliases
mail.byaddr		
publickey.byname		/etc/publickey
netid.byname		/etc/passwd, /etc/group, /etc/hosts, /etc/netid
netmasks.byaddr		/etc/netmasks
ypservers		(obtained from network broadcast)

Some NIS maps replace local AIX files so that when NIS is running, the information in the NIS maps is used instead of the information in the AIX files. Other NIS maps can be logically appended to local AIX files. This allows private and local information to be held on the NIS client.

The NIS master server is configured by running the `ypinit -m` command, which interactively asks questions and then builds the default maps by using the information supplied in the file /var/yp/Makefile. The NIS server can be started manually by using the command `startsrc -g yp`, or each daemon can be started individually by using `startsrc -s ypserv` (as an example). To configure the server daemons to start automatically, edit the file /etc/rc.nfs, or use the SMIT fastpath `smitty mkmaster`.

To rebuild the NIS maps:

```
# cd /var/yp

# make

couldn't find /etc/ethers
couldn't find /etc/networks
couldn't find /etc/netgroup
couldn't find /etc/bootparams
updated netid

pushed netid
couldn't find /etc/netmasks
Target "all" is up to date.
```

1.9.1.2 NIS Domains

DNS domains and NIS domains are not the same! DNS domains are parts of a *name space* used for resolving hostnames to IP addresses (like a tree structured directory, but for host names). An NIS domain is a set of information which is used to help administer a number of systems which are said to be in that NIS domain.

An NIS map is a mini-database which is built from AIX files and other available information. An NIS domain is a collection of NIS maps which will be used by one or more client systems. A client can belong to a number of NIS domains. You could, for example, have a domain which supplies hostname IP address mapping to all NIS client systems, but have two domains for user, password, and group information.

It is more common to have all clients in a department or on a LAN belonging to a single domain.

By default, NIS attempts to control hostname to IP address mapping. This conflicts with DNS, which also tries to do the same thing. NIS can be set up to operate with DNS so that any hostname resolution is first tried using DNS, and if the hostname is not found, the request is passed on to NIS. This assumes that the client and server are both set up to use DNS. Note that this is the opposite of the NIS implementations described in *Managing NFS and NIS, Hal Stern, O'Reilly 1994*. These non-AIX implementations will try NIS first, then DNS. If you want the AIX implementation of NIS to resolve names using NIS first, and then DNS, try the `-d` option when starting the `ypserv` daemon.

To enable NIS and DNS to interoperate, use an editor to change the /var/yp/Makefile file as follows:

- Locate the hosts.time stanza in the /var/yp/Makefile file.
- Change the two lines containing the word MAKEDBM:

  ```
  ...
  | $(MAKEDBM) - $(YPDBDIR)/$(DOM)/hosts.byname; \
  ...
  | $(MAKEDBM) - $(YPDBDIR)/$(DOM)/hosts.byaddr; \
  ...
  ```

  so that they look like:

  ```
  ...
  | $(MAKEDBM) -b - $(YPDBDIR)/$(DOM)/hosts.byname; \
  ...
  | $(MAKEDBM) -b - $(YPDBDIR)/$(DOM)/hosts.byaddr; \
  ...
  ```

1.9.1.3 NIS Clients

A client system can request information from any server that matches its domain (or domains). When the `ypbind` daemon starts, it broadcasts a request for each domain used by that client. The NIS server chosen continues to be used for that domain by the NIS client until the server is unavailable (the NIS timeout is reached). The NIS timeout defaults to 20 seconds. If you want a shorter timeout, you can set the environment variable YPTIMEOUT to the number of seconds to wait until timeout. The binding to the NIS server is then broken, and another request for an NIS server for that domain is broadcast.

Remember that an NIS server can also be a client of itself by running the ypbind daemon. This is usually the case; the NIS server runs both the server and client NIS daemons. In this document, NIS server refers to a server running the ypserv NIS server daemon, and NIS client refers to a system running the ypbind NIS client daemon, but not the ypserv daemon,

On the NIS client, you configure the system by starting the ypbind daemon and modifying /etc/passwd and /etc/group so that system routines reading these files will also use NIS information. This is done by using a plus sign in these files.

```
# cat /etc/passwd

root:!:0:0::/:/bin/ksh
daemon:!:1:1::/etc:
bin:!:2:2::/bin:
sys:!:3:3::/usr/sys:
adm:!:4:4::/usr/adm:
uucp:!:5:5::/usr/lib/uucp:
guest:!:100:100::/usr/guest:
nobody:!:4294967294:4294967294::/:
lpd:!:104:9::/:
+

# cat /etc/group

system:!:0:root
staff:!:1:davo,shoneen
bin:!:2:root,bin
sys:!:3:root,bin,sys
adm:!:4:bin,adm
uucp:!:5:uucp
mail:!:6:
security:!:7:root
cron:!:8:root
printq:!:9:lpd
audit:!:10:root
ecs:!:28:
nobody:!:4294967294:nobody
+:*:*
```

In the book *Managing NFS and NIS, Hal Stern, O'Reilly 1994*, it is recommended that for security reasons, you change the + in /etc/passwd to a line which reads

+:*:0:0:::

We found that with AIX V3.2.5 and AIX V4.1.4, this caused a problem with NIS, and we had to use a line without anything following the + to make NIS work correctly. The NIS problem was that NIS managed users could not log in to the NIS client. The line seen in the complete /etc/passwd above is the only line that would make NIS work correctly.

If the ypbind daemon is running, then the system is considered an NIS client. If you have set up the system to append NIS maps to /etc/passwd or /etc/group and ypbind is running, an NIS server must be available, or you will not be able to log in.

1.9.1.4 NIS User Names and Security

NIS users can log into any NIS client system which is in the NIS domain and configured to use the NIS passwd map. They can also use their own NIS password.

If a user changes his password on an NIS client system by using the yppasswd command, NIS will change the user's password in the NIS passwd.byname and passwd.byuid maps on the server. The changed NIS password will then be propagated to other slave servers. It will then be available to NIS client systems. When the yppasswd command is used on a client system, the NIS passwd map and password information in the /etc/security/passwd file are changed on the server.

When the passwd command is used on a client system for a user controlled by NIS, the yppasswd command is automatically invoked. This is true for both AIX V3.2.5 and AIX V4.1.4.

Attention!

In AIX V4.1.4, using either command, yppasswd or passwd, on an NIS master to change an NIS user password will not update the NIS passwd map.

The NIS passwd map will have to be rebuilt manually.

Note: With both AIX V3.2.5 and AIX V4.1.4, you might have a problem with the passwd map if the source /etc/passwd file contains blank lines.

The information in the NIS passwd.byname and passwd.byuid maps is logically appended to the information in the local /etc/passwd file. This means that it is possible to have users, such as root, who are administered locally on an NIS client using the normal AIX procedures. These non-NIS users on the NIS client can only log in on the local system, and have a password specific to the local system.

It is possible to change another user's password on a client system or on the master server. For example, the root user can change on the NIS master server any NIS user's password by using the yppasswd command. You will be asked to enter the old NIS password for this user. You can instead enter the root password. In AIX V3.2.5, once you have entered the user's new password and confirmed it, the NIS passwd map will automatically be rebuilt. In AIX V4.1.4, you must manually rebuild the passwd map.

As user root, you can also change on the NIS master any NIS user's password with the passwd command. In AIX V3.2.5 and AIX V4.1.4, you will have to rebuild the passwd map.

On any other system other than the NIS master server, even if you are logged in as user root, you must know the user's old password before you can change the password for that user. Table 3 on page 43 shows behavior of the passwd and the yppasswd commands in AIX V3.2.5 and in AIX V4.1.4, depending on which system the command is issued (NIS client or NIS server).

Table 3. Changing NIS User's Password as User root. Effects on NIS Maps and /etc/security/passwd.

AIX Version (Client and Server)	Command on NIS Server	Command on NIS Client	/etc/security/passwd updated?	NIS passwd map updated?	Notes
4.1.4	yppasswd	—	Yes— on server	No	Must rebuild NIS maps manually. Can use root's password or user's password when asked for Old NIS password.
	passwd	—	Yes— on server	No	Must rebuild NIS maps manually. Root is not prompted for old password.
	—	yppasswd	Yes— on server	Yes	root on NIS client must know user's old password.
	—	passwd	Yes— on server	Yes	root on NIS client must know user's old password.
3.2.5	yppasswd	—	Yes— on server	Yes	Required >5 character password. Can use root's password or user's old password when prompted for old NIS password.
	passwd	—	Yes— on server	No	Must rebuild NIS maps manually. Root is prompted for root password or user's old password.
	—	yppasswd	Yes— on server	Yes	Required >5 character password. Must use user's old NIS password.
	—	passwd	Yes— on server	Yes	passwd command uses yppasswd for NIS users. Must use user's old NIS password. Required >5 character password.

Note: If a user exists as both an NIS-administered user on the NIS master server and as a local user on an NIS client, then you should be careful which command is used to change the password for this user. This normally applies to the root user who is in the NIS passwd map generated from the /etc/passwd file on the NIS master server, and it may also be a local (non-NIS) user on the NIS client system. In this case, if the yppasswd command is used on the NIS client system, the NIS passwd map will be updated, and /etc/security/passwd will be changed on the NIS master server. If the passwd command is used on the NIS client system, /etc/security/passwd is only changed on the NIS client system, and the NIS passwd map is *not* updated.

The NIS passwd map can be rebuilt by using the following commands:

```
# cd /var/yp
# make passwd
```

Attention!

Once NIS is enabled on the server, many AIX security features can no longer be used. Many features which rely on user information in the files in the /etc/security directory are disabled because NIS was designed to only distribute basic UNIX security information.

NIS security can be improved by using the file /var/yp/securenets to restrict access to those clients with addresses within the IP range specified. NIS can also be run in secure mode. In secure mode, all clients and servers run the keyserv daemon. Public and private keys are used to make the transmission of passwords across the network more secure.

On the master server running AIX V4.1.4, you can designate local users to be administered by NIS rather than by AIX. The following section of /etc/security/user shows user *davo* who is administered by NIS, and user *stevo* who is administered by AIX.

```
stevo:
        admin = false

davo:
        admin = false
        registry = NIS
```

The registry information for a user can be changed by using `smitty chuser`. Set the registry to *NIS* for NIS-administered users and to *files* for AIX administration. If the registry line does not exist for a user, the default is to use AIX files as the default registry. If you blank-out the registry field in the SMIT panel, this will delete the registry line from the user stanza, and the registry will default to files.

Attention!

If you set a user to use NIS as their registry, then root user on the master server can no longer change the user's password without knowing the user's old password.

To reset the user's password, their registry should be changed to files. The root user on the NIS master server can then change the password, and the user's registry can then be set back to NIS.

1.9.1.5 NIS Netgroups

In order to assist in managing groups of users and hosts, NIS supports *netgroups*. Netgroups are definitions which consist of defined users, hosts and domains.

Netgroup entries have the form:

```
netgroup_name (host_name,user_name,domain_name) (host_name,.....)
```

Netgroups are used as a shorthand way to include host, user and domain information into NIS maps. Each NIS map only uses the information applicable to it.

We could define a netgroup called my_mates which includes users davo and shoneen and the systems drongo and derro.

```
my_mates (drongo,davo,labx) (derro,shoneen,labx)
```

In this example, if a system file or NIS map includes the netgroup my_mates, you would think that user davo is trusted only on system drongo and that user shoneen is only trusted on system derro. What really happens is that any system function that uses NIS for host information would trust hosts drongo and derro, and any system function that uses NIS for user information would trust davo and shoneen. For this reason, netgroups usually only include one type of information, either for hosts or users.

NIS netgroups can be used by the r commands and can be recognized in system files such as /etc/passwd and /etc/group.

For more information about administering netgroups, refer to *Managing NFS and NIS, Hal Stern, O'Reilly 1994*.

1.9.1.6 NIS Daemons

NIS relies on daemons on both client and server systems. These daemons are started from the /etc/rc.nfs script which is run at system startup. To change the startup options for the daemons, use the following table.

Table 4. NIS Daemons. Purpose and Configuration

NIS Daemon	Server or Client	Purpose	Configuration
ypbind	Client (server can also be a client)	Used to obtain information from NIS server	smitty mkclient or smitty rmypclient (can also edit /etc/rc.nfs)
ypserv	NIS master or slave server	Provides map information to clients	smitty chmaster or smitty rmypserv on the master server smitty chslave or smitty rmypserv on the slave server, (can also edit /etc/rc.nfs)
ypupdated	NIS master server	Prompts slave servers to update their maps	smitty chmaster or smitty rmypserv on the master server (can also edit /etc/rc.nfs)
yppasswdd	NIS master server	Processes requests to change user's passwords	smitty chmaster or smitty rmypserv on the master server (can also edit /etc/rc.nfs)
keyserv	NIS clients and servers running secure NIS	Provides security services for Secure NIS	smitty mkkeyserv or smitty rmkeyserv on the master server (can also edit /etc/rc.nfs)

The daemons can also be started individually. For example, to start the ypbind daemon, use the command:

```
# startsrc -s ypbind
```

On an NIS master server, you can start the ypbind, ypserv, ypupdated, and yppasswdd daemons by using:

```
# startsrc -g yp
```

1.9.2 NIS Migration Environment

The migration environment used to test NIS migration was:

- NIS master server- AIX V3.2.5 running NFS and acting as a DNS primary name server (NIS make file modified with -b flag). The system was also running as an NIS client of itself (ypbind daemon was running). The system was migrated to AIX V4.1.4 and checked with an AIX V3.2.5 client for proper operation.
- NIS client- AIX V3.2.5 running NFS and acting as a DNS client. The /etc/passwd and /etc/group files were modified to append NIS information. The AIX V3.2.5 NIS client was checked for correct operation with an AIX V3.2.5 NIS master server and then with the same NIS server running AIX V4.1.4. The NIS client was migrated to AIX V4.1.4 and checked with an AIX V3.2.5 NIS server and an AIX V4.1.4 NIS server for proper operation.

1.9.3 NIS Migration Planning

Attention!

Due to problems in the Migration Install logic, you must follow the procedures in 1.8.3, “NFS Differences and Migration Experiences” on page 31 to successfully complete NIS migration.

Attention!

The NIS ypbind daemon must be disabled in /etc/rc.nfs before migrating an NIS server system to AIX V4.1.4, or you will not be able to log in.

It is also recommended that ypbind is temporarily disabled on NIS client systems before migration. This allows the root user to log in and manually test NIS without being locked out by an NIS or TCP/IP problem.

The AIX V4.1.4 Migration Install process for the NIS and NFS filesets has a problem and does not replace some NIS and NFS files with the correct AIX V4.1.4 default files. In particular, the lines in /etc/rc.nfs that start the NIS server will not work correctly. This is because the symbolic link from /etc/yp to /var/yp has been removed during migration. This causes a test for the existence of an NIS map directory to fail, and the ypserv daemon is not started. Since the server process does not start, any system, client *or server*, that runs the ypbind daemon automatically at boot time will not be able to contact an NIS server and will not allow you to log in.

If you are reading this before doing a Migration Install of AIX V4.1.4, you should edit the file /etc/rc.nfs, and change

```
...
if &lbr. -x /usr/etc/ypbind &rbr.; then
      startsrc -s ypbind
fi
...
```

to

```
...
# if &lbr. -x /usr/etc/ypbind &rbr.; then
#       startsrc -s ypbind
# fi
...
```

After migration, follow the steps in 1.8.3, “NFS Differences and Migration Experiences” on page 31 to recover the AIX V4.1.4 default files and merge them with the AIX V3.2.5 customized files.

If you are installing AIX V4.1 by using the Preservation or Overwrite Install methods, then you should note which NIS daemons are normally started on the system you are migrating. You should also note the domain name to which the system belongs.

Some files which you may consider backing up on your master server:

- /var/yp/Makefile- used to create NIS maps. Often modified to allow DNS and NIS to coexist.
- /var/yp/updaters- used to create NIS maps for running secure NFS.

You should also check Table 2 on page 38 for other files which are used to build NIS maps on your master server. These files should be backed up. You should also use the `ypcat -k map_name` command on the NIS master server to document the information in your NIS maps. Remember to do all of your domains if your master supports more than one domain!

1.9.4 NIS Differences and Migration Experiences

Attention!

A problem with AIX V4.1.4 migration means that some NFS and NIS files must be manually recovered from the AIX distribution tape or CD and manually merged with the customized AIX V3.2.5 files left on the system. If you are running NIS, you should do a recovery and merge of all of the NFS and NIS files. See 1.8.3, “NFS Differences and Migration Experiences” on page 31 for further information.

After migration of an NIS server system which was also running the `ypbind` daemon, it was not possible to log in to the system. The reason for this is documented in 1.9.3, “NIS Migration Planning” on page 46. If you want to recover from this:

1. Reboot the system in Service mode from either an AIX V4.1.4 distribution tape/CD or an AIX V4.1.4 `mksysb` tape.
2. Choose the **System Maintenance** option, then the option to **Access the rootvg Volume Group**. Choose the appropriate disk to access, then the option to **Start a Shell**.
3. `export TERM=lft` (or whatever type of terminal you are using- ibm3151, vt100, ansi, wy60, and so on).
4. `vi /etc/rc.nfs`. Comment out the lines starting the `ypbind` daemon by inserting a # at the start of the line. Exit `vi`.
5. Put the system key to Normal mode.
6. `shutdown -Fr`

7. Once you have rebooted, perform the recovery/merge of the NFS and NIS files as documented in 1.8.3, “NFS Differences and Migration Experiences” on page 31.

When setting up NIS on the client (running AIX V3.2.5), there were problems when the line in /etc/passwd containing the NIS token + contained any other information. The dummy password and group information on the NIS token line, as recommended in *Managing NFS and NIS, Hal Stern, O'Reilly 1994*, (+:*:0:0::::), caused the NIS client to ignore any NIS user information. It seems that the longer line is not recognized as an NIS token.

After the client was upgraded to AIX V4.1.4, NIS was found to still operate correctly. The problem with /etc/passwd containing an NIS token line with more than + in it was still present in the NIS client running AIX V4.1.4.

The problem with building NIS passwd maps from an /etc/passwd with blank lines in it was noted in both AIX V3.2.5 and AIX V4.1.4.

The symbolic link from /etc/yp to /var/yp is no longer present after migration to AIX V4.1.4.

The `ypbind` daemon has different default behavior in AIX V4.1.4. The default behavior is now to reject all requests from any `ypset` command. New flags for the `ypbind` daemon are:

`-s` to run `ypbind` in secure mode. The daemon will use only privileged ports for communication.

`-ypset` to allow `ypbind` to accept `ypset` commands from the local host and remote hosts.

`-ypsetme` to make `ypbind` only accept `ypset` commands from the local host.

In AIX V3.2.5, the NIS password had to be a minimum of six characters in length. In AIX V4.1.4, this restriction does not apply; passwords can be shorter.

When a root user logged in on the NIM master server changes a user's password by using the `yppasswd` command:

- In AIX V3.2.5, both the NIS passwd map and /etc/security/passwd on the NIS master server are updated.
- In AIX V4.1.4, /etc/security/passwd is updated on the NIS master server. The NIS passwd map is not automatically updated; so you must rebuild the NIS map manually (see instructions on page 43).

1.9.5 NIS Interoperability - AIX V3.2.5 and AIX V4.1.4

NIS worked properly between AIX V3.2.5 clients and an AIX V4.1.4 server. No problems were encountered.

1.10 Migrating ftp

If you want to easily move files between systems, the use of the File Transfer Protocol (ftp) is one way to do it. The ftp file transfer method is supported on almost any computer system which supports TCP/IP networking.

Using ftp to another system, you are asked to log in with a valid user name and password before being allowed to transfer files.

On the server system, the ftpd daemon is started by the inetd daemon as soon as an ftp service is requested. By default, ftp is enabled on AIX and can be disabled by commenting out the ftpd line in /etc/inetd.conf. The inetd daemon must then be informed that the configuration file has changed by using the command:

```
refresh -s inetd
```

If the ftp server system has been set up as an *anonymous ftp* server, then users can supply a user name of ftp or a username of anonymous, and then supply an E-mail address as the password. This is often used on the Internet to provide global access to publicly available software.

Setting up an anonymous ftp server really just means adding a special user and creating a special directory structure. The system continues to run the ftpd daemon as normal. If you log in to an ftp server as ftp or anonymous, you are restricted to access files in the anonymous ftp directory and its subdirectories. This is done by using the chroot command when the ftp/anonymous user logs in. The anonymous ftp user cannot cd to any directory outside the directory defined. For example, if the home directory of the ftp user is /home/my_ftp_dir, then anonymous ftp users see /home/my_ftp_dir as /, and they cannot cd to /home or any other directory above /home/my_ftp_dir.

If you want to easily configure an anonymous ftp server, use the shell script /usr/samples/tcpip/anon.ftp (AIX V4.1.4) or /usr/lpp/tcpip/samples/anon.ftp (AIX V3.2.5). This sets up the ftp and anonymous users as well as the proper directory structure, permissions and contents. The shell script uses the ftp user's home directory as the *root directory* for anonymous ftp users. If you want to use a directory other than /home/ftp, you can create a user called ftp with a different home directory, for example /home/my_ftp_dir, and then run the anon.ftp shell script. The shell script will give a warning that the ftp user already exists, but will complete successfully.

Attention!

When using the anon.ftp script to set up an anonymous ftp server on an NIS client, you should stop the NIS ypbind daemon before executing this script. Restart the ypbind daemon after the script has completed successfully.

1.10.1 ftp Migration Environment

Both systems were tested for correct ftp operation with AIX V3.2.5. One system (yobbo) was set up as a DNS client and an NIS client. One system (hoon) was set up as a DNS name server, an NIS master server and an anonymous ftp server.

The anonymous ftp server was migrated to AIX V4.1.4 and tested for correct operation for both normal ftp and anonymous ftp. This was done using the client running AIX V3.2.5. No problems were found for normal ftp users. Some problems were found using this system as an anonymous ftp server after migration.

The client system was then migrated to AIX V4.1.4 and tested for correct operation with the server. This was done for both normal ftp and anonymous ftp. No problems were found for normal ftp or anonymous ftp users.

1.10.2 ftp Migration Planning

There were no special requirements for planning the migration of ftp services to AIX V4.1.4 when using the Migration Install method.

1.10.3 ftp Differences and Migration Experiences

After migrating the system to AIX V4.1.4, the `ls -l` command would not work for anonymous ftp users although the `ls` command continued to work.

The reason for this is that anonymous ftp users are placed into their own restricted set of directories by using the chroot command as part of their login sequence. The home directory of the user called ftp becomes the root directory for anonymous ftp users. In the examples which follow, the ftp user was defined with a home directory of /home/my_ftp_dir.

When the shell script anon.ftp is executed to set up the system to support anonymous ftp, a copy of the `ls` command is placed in /home/my_ftp_dir/bin (in this example), and a copy of the libc.a library is placed in /home/my_ftp_dir/lib. This is done so that the anonymous ftp user can find and execute the `ls` command.

If the anon.ftp shell script is executed to set up anonymous ftp on a system running AIX V3.2.5, and this system is then migrated to AIX V4.1.4. The `ls` and the library file libc.a found by the anonymous user in /home/my_ftp_dir/bin and /home/my_ftp_dir/lib will still be at the AIX V3.2.5 level.

In order to correct this, follow the steps below after migration to AIX V4.1.4 has completed. These steps copy the AIX V4.1.4 ls and libc.a into the directory tree used by anonymous ftp users.

```
# cp /usr/bin/ls /home/my_ftp_dir/bin

# cp /usr/ccs/lib/libc.a /home/my_ftp_dir/lib
```

1.10.4 ftp Interoperability - AIX V3.2.5 and AIX V4.1.4

There were no problems found in interoperability between AIX V3.2.5 and AIX V4.1.4 ftp services.

Part 2. SNA

Chapter 2. SNA Migration

Migrating a system with AIX Version 3.2 and a Systems Network Architecture (SNA) product installed to AIX Version 4.1 is not too complex; however, there are some AIX Version 4.1 specifics that you should consider. These include changes to the packaging of some AIX components and changes to the format and behavior of some commands.

If you are running SNA Server/6000 Version 2.1 on your AIX Version 3.2 system, the migration is easy because your existing SNA profiles can be used after the migration to Communications Server Version 4 for AIX running on AIX Version 4.1. However, if you are currently using SNA Services/6000 Version 1.2.1, you will have to perform the additional step of migrating your SNA profiles.

This chapter explains the migration process and notes some details that relate specifically to the implementation of SNA on AIX V4.1, including:

- Migration methodology
- Installing Communications Server Version 4 for AIX
- Installing HCON Version 2.1
- Migrating SNA profiles
- Migrating SNA API programs
- Saving and restoring HCON profiles
- A migration example

2.1 Related Publications

The following books are related directly to SNA and other AIX-SNA products. For books related to AIX in general, and AIX migration, see “Related Publications” on page xviii.

- *AIX Communications Server: Up and Running*, SC31-8247
- *SNA Server for AIX General Information Manual*, SC31-8198
- *SNA Server for AIX User's Guide*, SC31-8211
- *SNA Server for AIX Configuration Reference*, SC31-8213
- *SNA Server for AIX Command Reference*, SC31-8214
- *SNA Server for AIX Diagnosis Guide and Messages*, SC31-8215
- *APPC Application Suite for AIX User's Guide*, SC31-8218
- *SNA Server for AIX Transaction Program Reference*, SC31-8212
- *CPI Communications for SNA Server*, GC31-8210
- *SNA Server for AIX Gateway User's Guide*, SC31-8216
- *SNA Server for AIX Planning and Performance Guide*, SC31-8094
- *SNA Server for AIX AnyNet Guide to Sockets over SNA*, SC31-8217
- *SNA Server for AIX AnyNet Guide to APPC over TCP/IP*, SC31-8221
- *SNA Server for AIX Channel Connectivity User's Guide*, SC31-8219

- *RISC/6000 to Mainframe Using S/370 Channel Connections*, SG24-4589.
- *AIX Version 4.1 Enterprise Systems Connection Adapter: User's Guide and Service Information*, SC31-8197
- *AIX Version 4.1 Block Multiplexer Channel Adapter: User's Guide and Service Information*, SC31-8196
- *AIXLink/X.25 1.1 for AIX: Guide and Reference*, SC23-2520
- *AIX/6000 X.25 LPP Cookbook*, GG24-4475

Softcopy Publications: The information in the SNA Server hardcopy manuals is also available in a DynaText database for the RISC System/6000 workstation. You can use this library to search for specific information or to view online versions of each of the SNA Server books.

To use and view this book, you have to install the browser and the publication filesets which are shipped with the Communications Server Version 4 product. DynaText is available only in the AIXwindows environment.

The browser can be started by the command `dtext` or by using `smit snapubs`.

2.2 Preparing for SNA Migration

The information in this chapter assumes that you have carefully read the sections on preparing your system for the migration to AIX Version 4.1. Preparing for migration of SNA is not very different, but some aspects should be taken into account.

The following list will give you some ideas of topics you have to think about.

- What level of SNA is installed?
  - SNA Services/6000 Version 1.2.1 or SNA Server/6000 Version 2.1
- What functions do I use?
  - Only SNA basic functions such as Logical Unit (LU) 1, 2, 3 and LU 6.2
  - SNA LU type 0 functions
  - SNA Gateway function
  - AnyNet function
  - Channel Support
  - Advanced Peer-to-Peer Networking (APPN) function
- What products must I order from IBM?
  - Depends on the functions above
- What other programs are using SNA—do they need an update too?
  - Host Connection Program/6000 (HCON) or other 3270 emulator
  - Remote Job Entry (RJE) emulation
  - LU 6.2 APPC programs
  - NetView Service Point or other systems management products
  - Other products or programs
  - Are all product delivery media available for re-install?

- What scripts or other procedures are written for SNA?
  - Any special scripts for controlling SNA or SNA-dependent products
  - Procedures for restart and recovery
- What hardware do I use today for SNA communication?
  - Fiber Distributed Data Interface (FDDI), X.25 Coprocessor/2, X.25 Portmaster/A,
  - 4-port Multi-Protocol Adapter, Single-Port Multi-Protocol Adapter
  - Token-Ring, Ethernet
  - Enterprise System Connection (ESCON) channel, Block Multiplexer (BLKMUX) channel
- Is that hardware supported (possibly on a new system)?
  - Check if these adapters are available on Industry Standard Architecture (ISA) or Peripheral Component Interconnect (PCI) bus machines.
  - Do I need special driver software, new versions or new programs?
- How much time do I need for migration?
  - Depends on used functions and programs
  - Is a test system available for testing some functions in advance?
- What happens when the migration does not work?
  - Is there any way back?

There may be other considerations, but these are the major questions you must answer before you should start the migration of SNA. The next chapters give you more detailed information on most of these questions. Some of these questions are answered in *A Holistic Approach to AIX V4.1 Migration, Volume 1*, SG24-4652, and general hardware and software availability can be found in *A Holistic Approach to AIX4.1 Migration, Planning Guide*, SG24-4651. However, some questions depend on your specific system environment. There may be SNA applications from other vendors or scripts and programs written for the maintenance of your SNA environment. It is not possible to cover all possible questions and cases in these books; however, the examples present should help you to understand the implications of the migration.

2.3 History of SNA and HCON

Over the life of AIX Version 3.2 and Version 4.1, the SNA product set has undergone continual evolution. The levels of SNA products and their associated AIX levels are shown in Table 5 on page 56. To determine the level of SNA that is currently installed on your system, use the command:

```
$ lslpp -L "*sna*" "*gw*"
```

By comparing the output with the table, you can determine the version of SNA currently installed.

Table 5. SNA Product History

AIX Req.	SNA Product	Date Announced	Base Object	lslpp output
3.2	IBM AIX System Network Architecture Services/6000 Version 1.2	Jan. 21, 1992	sna.sna.obj	1.2.0.0 - 1.2.101.510
3.2.3	IBM AIX SNA Server/6000 Version 2.1	Dec. 21, 1993	sna.sna.obj	1.3.93.495 - 1.3.94.110
	IBM AIX SNA Gateway/6000 Version 2.1		gw.sna.obj	1.3.93.495 - 1.3.94.110
	IBM AIX SNA Server/6000 Version 2.1.1	July 26, 1994	sna.sna.obj	1.3.94.231 - 1.3.95.20
	IBM AIX SNA Gateway/6000 Version 2.1.1		gw.sna.obj	1.3.94.231 - 1.3.95.20
	AnyNet/6000 APPC over TCP/IP		anynet.snaip.obj	1.1.94.462 - 1.1.95.70
	AnyNet/6000 Sockets over SNA		anynet.snackets.obj	1.1.94.462 - 1.1.95.70
	IBM SNA Application Access for AIX Version 1.1	Oct. 4, 1994	SNA_AA.rte.obj	1.1.0.0
	IBM SNA Client Access for AIX V1.1 or Version 1.2		SNA_CA.rte.obj	1.1.0.0 - 1.2.0.0
	IBM Desktop SNA for AIX Version 1.1		snadt.rte.obj	1.1.94.32
	IBM AIX SNA Server/6000 Version 2.1.2	Feb. 28, 1995	sna.sna.obj	1.3.95.50 - 1.3.95.293
	IBM AIX SNA Gateway/6000 Version 2.1.2		gw.sna.obj	1.3.95.50 - 1.3.95.293
	SNA Channel Connectivity for ESCON		sna.escon.cuu and sna.escon.usr	1.3.95.55
	SNA Channel Connectivity for Block Multiplexer		sna.blkmux.cuu and sna.blkmux.usr	1.3.95.55
4.1.1	IBM AIX SNA Server/6000 Version 2.2	March 21, 1995	sna.rte	2.2.0.1
	IBM AIX SNA Gateway/6000 Version 2.2		gw.rte	2.2.0.1
4.1.2	IBM AIX SNA Server Version 3.1	Sep. 12, 1995	sna.rte	3.1.0.0
4.1.2	IBM AIX Communications Server Version 4	March 12, 1996	sna.rte	3.1.1.0

Similarly, there have been several versions of HCON over the life of AIX V3.2.5 and V4.1. Table 6 on page 57 shows the relationship between AIX and HCON levels. To determine the level of HCON that is currently installed on your system, use the command:

```
$ lslpp -L "*hcon*"
```

Table 6. HCON Product History					
AIX Req.	**SNA Req.**	**HCON Version**	**Date Announced**	**Base Object**	**lslpp output**
3.2	1.2	IBM AIX 3270 Host Connection Program/6000 Version 1.3	Jan. 21, 1992	hcon.obj	1.3.0.0
3.2.3E	1.2	IBM AIX 3270 Host Connection Program/6000 Version 1.3.1	Sep. 22, 1992	hcon.obj	1.3.0.0
3.2.5	1.2 or 2.1	IBM AIX 3270 Host Connection Program/6000 Version 1.3.2	May 24, 1994	hcon.obj	1.3.0.0
4.1.1 4.1.2	2.2 3.1 or 4	IBM AIX 3270 Host Connection Program for AIX Version 2.1	Oct. 4, 1994	hcon.rte	2.1.0.0 - 2.1.1.0

The table shows several different SNA products. The differences between these products is explained in the following sections. Our testing was performed using the following products:

- SNA Services/6000 Version 1.2.1
- SNA Server/6000 Version 2.1
- SNA Gateway/6000 Version 2.1
- Communications Server Version 4 for AIX
- HCON Version 1.3
- HCON Version 2.1

The other products and versions listed and described below are for information and positioning purposes only.

2.3.1.1 SNA Services/6000 Version 1.2

SNA Services/6000 was the first SNA product available on AIX Version 3. It provided communications with applications running on other systems supporting the SNA protocol. This version is no longer marketed or supplied and was only be supported until December 31, 1995. It is superseded by SNA Server/6000.

SNA Services/6000 V1.2 had one modification level and some minor enhancement Program Temporary Fixes (PTFs). The final version, SNA Services/6000 V1.2.1, included these functions:

- Type 2.1 nodes
- Application Programming Interface (API) for Advanced Program-to-Program Communications (APPC) including Common Programming Interface for Communications (CPI-C)
- Multiple LU support for a maximum of 5000 sessions per node
- Support for the following link types:
  - Synchronous Data Link Control (SDLC)
  - Token-Ring
  - Ethernet
  - X.25

When compared to later SNA products, SNA Services/6000 had the following limitations:

- No Advanced Peer-to-Peer Networking (APPN) support, but can be connected as a Low Entry Networking (LEN) node to an APPN network

- No support for syncpoint
- Multiple Physical Units (PUs) not supported

2.3.1.2 SNA Server/6000 Version 2.1

This is the successor to SNA Services/6000 and is the current base SNA product for users on AIX Version 3.2. This product includes these functions:

- Full APPN function for the following type 2.1 nodes.
  - Low Entry Networking (LEN) node
  - End Node (EN)
  - Network Node (NN)
- API for APPC including Common Programming Interface for Communications (CPI-C)
- Syncpoint support
- Multiple PU support
- Extended LU 0 function

This version has two modification levels.

SNA Server/6000 Version 2.1.1: These functions are added or extended by the 2.1.1 modification level:

- FDDI adapter (needs AIX V3.2.5 or above)
- Multiple PUs for SDLC lines (need AIX V3.2.5 or above and PTF U422520)
- APPC session timeout
- AnyNet/6000 APPC over TCP/IP support (separately orderable)
- AnyNet/6000 Sockets over SNA support (separately orderable)

SNA Server/6000 Version 2.1.2: This modification level requires AIX V3.2.5 or above. These functions are added or extended by this modification level:

- Block Multiplexer SNA channel connection support
- ESCON SNA channel connection support
- SNA Application Suite for AIX (APPC application toolkit)
- APPC U-shaped session support (for Transaction Program testing with only one local node)
- Improved performance

2.3.1.3 SNA Gateway/6000 Version 2.1

SNA Gateway/6000 provides a gateway function for connecting between a host and multiple downstream workstations. It supports the connection of clients over X.25, SDLC, token-ring, and Ethernet.

SNA Gateway/6000 Version 2.1 requires that Version 2.1 of SNA Server/6000 be installed. SNA Gateway also has two modification levels. In each case, it is important to ensure that the levels of SNA Server/6000 and SNA Gateway/6000 remain in synchronization.

SNA Gateway/6000 Version 2.1.1: The following functions are added or extended by the 2.1.1 modification level.

- FDDI adapter support
- Multiple PU for SDLC lines (not for primary station)
- Quick configuration menu
- Improved performance

SNA Gateway/6000 Version 2.1.2: This modification level does not provide any new functions. It exists only to make required adjustments allowing SNA Gateway to operate with Version 2.1.2 of SNA Server.

2.3.1.4 SNA Server/6000 Version 2.2

This was the first SNA product for AIX Version 4.1 on uniprocessor systems only. This product is no longer marketed or supplied. Service ended on April 26, 1996. It is superseded by SNA Server Version 3.1 for AIX or by Communications Server Version 4 for AIX. This product had most of the functions as SNA Server/6000 Version 2.1.2; however, the following functions were *not* supported:

- Block Multiplexer SNA channel connection
- ESCON SNA channel connection
- AnyNet/6000 APPC over TCP/IP

2.3.1.5 SNA Gateway/6000 Version 2.2

SNA Gateway/6000 Version 2.2 provided gateway functions for uniprocessor systems running AIX Version 4.1. The SNA Gateway/6000 Version 2.2 is withdrawn from marketing and service ends together with SNA Server/6000 Version 2.2 at the end of April 1996. The product is superseded by Communications Server Version 4 for AIX.

Version 2.2 of SNA Gateway/6000 required that Version 2.2 of SNA Server/6000 be installed on the system. The functions were identical to the SNA Gateway/6000 Version 2.1.1 on AIX V3.2.5 systems.

2.3.1.6 SNA Server Version 3.1 for AIX

This SNA Server Version 3.1 for AIX is the first version which supports both uniprocessor (UP) and multiprocessor (SMP) machines. The scope of functions supported can be compared to all functions available in AIX V3.2.5 systems running SNA Server/6000 Version 2.1.2 but with some minor enhancements, like the Simple Network Management Protocol (SNMP) agent for APPN. The main difference is the new packaging of the products:

- SNA Server Version 3.1 (5765-582) includes the products:
  - SNA Gateway for AIX
  - AnyNet Sockets over SNA
  - AnyNet APPC over TCP/IP
- The following options are now moved into separate products:
  - ESCON Channel Connectivity for AIX V1.1, (5765-603)
  - BLKMUX Channel Connectivity for AIX V1.1, (5765-604)

The SNA Server Version 3.1 product comes with all functions, softcopy publications and tools included on the delivery medium. It is up to the user to select the required functions during the installation process.

SNA Server Version 3.1 requires AIX Version 4.1.2 or higher to be installed on the system.

Since this version of SNA has so many differences from previous versions in ordering and handling, a more complete description is given in 2.4, “SNA Product Ordering” on page 64.

Note: The product is superseded by Communications Server Version 4 for AIX. From the migration stand point, there are no differences between migrating to SNA Server Version 3.1 and migrating to Communications Server Version 4

2.3.1.7 Communications Server Version 4 for AIX

SNA Server Version 3.1 for AIX has been enhanced and incorporated into the Communications Server Version 4 for AIX (5765-652) and is part of the new IBM Software Server family of offerings.

The new Communications Server Version 4 for AIX includes all the characteristics of the existing SNA Server Version 3.1 product along with the following enhancements:

- Common installation procedures with other Software Servers for AIX
- Bundled 3270 Host Connection emulation capability (single session only)
- Bundled AnyNet: APPC over TCP/IP Gateway
- API enhancement (APPC Non-Blocking)

Communications Server Version 4 requires AIX Version 4.1.2 or higher to be installed, but the new Software Server installation function requires AIX Version 4.1.4 to be installed on the system.

Since all functions and procedures, except those mentioned above in the enhancements, are fully compatible with SNA Server Version 3.1, the migration to Communications Server Version 4 is identical to SNA Server Version 3.1 migrations.

The 3270 Host Connection program (HCON) is bundled with Communications Server Version 4 and comes on the same media, but allows only a single session to be used. For the full multi-user usage, the HCON product 5765-398 must be ordered.

Ordering and pricing procedures are similar to SNA Server Version 3.1. For basic information, refer to 2.4, “SNA Product Ordering” on page 64. Detailed information about session counting and ordering is given in the announcement letter.

2.3.1.8 AnyNet/6000 APPC over TCP/IP

This product supports the execution of APPC and CPI-C applications over a TCP/IP network. It requires that SNA Server/6000 Version 2.1.1 or above be installed on the system. AnyNet/6000 APPC over TCP/IP is now part of the SNA Server Version 3.1 for AIX or Communications Server Version 4 for AIX.

2.3.1.9 AnyNet/6000 APPC over TCP/IP Gateway

This support was added with Communications Server Version 4. APPC over TCP/IP gateway for AIX connects Internet Protocol (IP) networks and SNA networks to enable communications between SNA applications. In conjunction with any other AnyNet product, SNA applications can now run on any IP-attached workstation or host to any SNA application on native SNA networks.

2.3.1.10 AnyNet/6000 Sockets over SNA

AnyNet/6000 Sockets over SNA allows applications written to use the Berkeley Software Distribution (BSD) sockets interface to communicate over an SNA network. It requires that SNA Server/6000 Version 2.1.1 or above be installed on the system. On AIX Version 4 systems, it is part of the SNA Server Version 3.1 for AIX or Communications Server Version 4 for AIX.

2.3.1.11 SNA Application Access for AIX Version 1.1

SNA Application Access provides a consistent platform for allowing SNA devices located anywhere throughout an SNA network to communicate with AIX applications residing on the RISC System/6000. These devices could include 3270 terminals, printers or other resources. SNA Application Access includes the appropriate level of function to provide the SNA communications capabilities of a host system and a communications controller, and it positions the RISC System/6000 as an SNA applications server.

SNA Application Access requires that AIX Version 3.2.3 or above and SNA Server/6000 be installed on the system. Currently, it is not supported on AIX Version 4.1.

2.3.1.12 SNA Client Access for AIX

SNA Client Access, in conjunction with AIX SNA Server/6000, provides a consistent platform for client programs to access an SNA network. It functions as a TCP/IP server, and it provides an SNA network access to client applications running anywhere in a TCP/IP network. Since SNA Client Access processes all the lower-level SNA protocols and provides service management, the client programs can focus on sharing information with the host systems applications. Client programs can attach to SNA Client Access to gain access to host system applications, such as JES2, TSO, POWER, IMS, CICS, and NetView. SNA Client Access Version 1.1 supports all of the clients required for tn3270 and tn5250 protocols, while SNA Client Access Version 1.2 supports, in addition, the tn3270e protocol, which allows 3270 printer emulation over a TCP/IP network.

AIX Version 3.2.3 or above and SNA Server/6000 are required to run SNA Client Access. SNA Client Access is supported on AIX Version 4.1. There are no additional fixes or enhancements required to provide the support.

2.3.1.13 IBM Desktop SNA for AIX Version 1.1

Desktop SNA is tailored for the single, end-user, desktop environment. It has reduced function from a network server perspective, but still provides many of the traditional SNA functions that are available in AIX SNA Server/6000 V2.1. The reduced disk space requirement of 10 MB, compared with the 23 MB required to run SNA Server/6000, is a major benefit for client systems.

Desktop SNA is supported on AIX Version 3.2.3 and above but is not supported on AIX Version 4. Compared to SNA Server/6000, it has the following limitations:

- A maximum or 100 active sessions

- No more than eight sessions of independent LUs per Link Station
- No support for SDLC or X.25 wide area networks
- No support for FDDI
- No LU type 0 API support
- Generic SNA API not supported
- No support for APPN network nodes

2.3.1.14 IBM 3270 Host Connection Program for AIX V1.3

IBM 3270 Host Connection Program for AIX (HCON) provides host connection with 3270 display sessions over SNA or TCP/IP protocols. In the SNA environment, HCON 1.3 will operate with either SNA Services/6000 or SNA Server/6000 on AIX Version 3.2.5.

The main functions of HCON are:

- 3278/3279 display terminal emulation
- 3286/3287 printer emulation
- DFT display session
- Extended data stream support
- File transfer to and from mainframe systems
- API for writing programs to communicate with a mainframe system
- Multiple host connections
- PU type 2.1 support (requires SNA)

HCON 1.3 has two modification levels.

IBM 3270 Host Connection Program for AIX V1.3.1: These functions are added or extended by this modification level:

- High-Level Language Application Programming Interface (HLLAPI) support
- 3287 printer emulation, including background printing
- SNA connection for 3172 LAN terminals over IEEE 802.3 Ethernet
- Channel connection to host with TCP/IP
- Korean and Chinese language support

IBM 3270 Host Connection Program for AIX V1.3.2: These functions are added or extended by this modification level:

- AIXwindows/Motif base interface support
- Light pen function
- Title bar of session control, utility and PF-key bar
- HLLAPI
- Wyse 370 - ASCII terminal support

2.3.1.15 IBM 3270 Host Connection Program for AIX V2.1

This product provides Host Connection with 3270 display sessions on SNA or TCP/IP for systems running AIX Version 4.1. It requires AIX Version 4.1.1 or above for use over TCP/IP networks and AIX Version 4.1.2 or above and SNA Server Version 3.1 or above for use over SNA networks.

These functions are added or are different from Version 1.3.2.

- Support for AIXwindows Desktop Manager
- Hebrew and Arabic Language support
- HCON API library name is changed from libg3270.a to libhllapi.a

A single-session version of the IBM 3270 Host Connection Program is now bundled with Communications Server Version 4; however, the full product must be purchased to provide multi-session support.

2.3.1.16 Channel Connectivity on AIX

Channel connectivity to /370 or /390 host systems was first available for TCP/IP connections only. To use channel connectivity, it was necessary to order the special adapter card and device driver for your AIX system. In February 1995, SNA Server/6000 2.1.2, which expands the channel connectivity to SNA traffic, was announced. The old device driver was replaced by a new one, which became part of the SNA Server product. This device now driver supports both SNA and TCP/IP traffic.

In September 1995, IBM announced SNA Server Version 3.1 for AIX Version 4.1 systems. Part of that announcement was two new program products—the channel connectivity products for ESCON and BLKMUX channels. The support for the different protocols, such as TCP/IP, CLIO/S and SNA, have been split up. The base product allows TCP/IP traffic only. In order to use other protocols, such as Client Input Output/Sockets (CLIO/S), a data transfer protocol for high-speed file and data transfer or SNA, separate programs and features are available.

BLKMUX Channel Connectivity for AIX V1.1: This product (5765-604) allows the connection of a RISC System/6000 to a mainframe using the Block Multiplexer channel adapter. Protocols supported are TCP/IP, CLIO/S, and SNA. To use these protocols, you need the adapter card and the following extensions:

TCP/IP No further features

CLIO/S CLIO/S licensed program (5648-129) Version 2.1

SNA SNA Server Version 3.1 or Communications Server Version 4 plus the SNA Channel Connectivity feature

Note: The SNA Channel Connectivity feature provides the *dlcchannel* device function which enables the access from SNA to the channel device drivers.

ESCON Channel Connectivity for AIX V1.1: This product (5765-603) allows the connection of a RISC System/6000 to a mainframe using the ESCON channel connection. Protocols supported are TCP/IP, CLIO/S and SNA. :i2=sna.ESCON support To use these protocols, you need the adapter card and the following extensions:

TCP/IP No further features

CLIO/S CLIO/S licensed program (5648-129) Version 2.1

SNA SNA Server Version 3.1 or Communications Server Version 4 plus the SNA Channel Connectivity feature

2.4 SNA Product Ordering

Before you install and migrate the SNA products, you must first order and obtain the new software. You must order SNA Server Version 3.1 for AIX, product number 5765-582, or Communications Server Version 4 for AIX, product number 5765-652. If you already have the SNA software (possibly a demo or trial license) and only require a full license, you should order only the SNA product. If you require a tape or CD shipment of software, you must also order an AIX System Program Order (SPO), product number 5692-AIX, specifying the apropriate SNA features.

The following sections will give you a basic understanding of the pricing structure of the SNA product in order to help you to place a valid order.

2.4.1 SNA Pricing Structure

The pricing structure of all previous versions or releases of SNA Server/6000 was easy to understand and use, but was sometimes not very fair. It was based on the processor group of the machine. For small machines, it was cheap; for big machines, it could be expensive. In general, the pricing did not depend on usage.

This has now changed. In AIX 4.1, many products are priced on usage. For Communications Server Version 4, the usage is counted by SNA *sessions* instead of by the number of users. This is your first problem—when you order the Communications Server Version 4, you have to know the number of SNA sessions. The different session types are listed below. You will need to consider the required system environment to determine the number of each type of session that will be used.

LU 1, 2 and 3 and 6.2-dependent LU-LU sessions

These session types are mainly used for 3270 emulation. When used in conjunction with HCON, you may simply count the HCON-users and multiply by the number of sessions for each user. Add every additional SNA printer session or dependent LU 6.2 session to get the total number of these types of sessions. You may also approach this question in a host environment by finding the total number in the Virtual Telecommunication Access Method (VTAM) or Network Control Program (NCP) definitions and using this as the total number of all LUs defined to the PU (Physical Unit) for this server. The controlling System Services Control Program-Logical Unit (SSCP-LU) session for each of the LUs is not counted, just the LU-LU session when connected to an application.

Independent LU 6.2 user sessions

Every SNA *conversation* must be counted. Since this number is hard to determine, most times you will find the right number by counting the number of *users* using this type of communication.

APPN intermediate sessions

In APPN networks, sessions from other nodes can go through this server. For example, Low Entry Networking (LEN) nodes must use this type of service.

SNA Gateway sessions

All dependent LU sessions for downstream-connected SNA devices have to be counted. But be careful—they count only once for the whole connection, from downstream up to the connected host. So, you may have already counted then in step 1 of these rules.

LU 0 sessions

Every LU 0 primary session has to be counted, but will not be checked. These types of sessions are the controlling sessions for connected LU 0 devices, like teller machines or cash points. Every LU 0 secondary session has to be counted and will be checked, with any use above the licensed level being listed in the SNA Server session log. These types of sessions are used for special applications operating under the LU 0 protocol. For example, the NetView Distribution Manager can be used with either an LU 6.2 connection or an LU 0 connection.

Note

The sessions used for Client Access for AIX are not counted on the SNA Server license. This product is licensed according to the number of users and includes the usage of SNA sessions.

In order to place a valid initial order, you must select at least one session. A valid initial order, therefore, consists of:

- Base product order
- Session order

The session order is divided into different Use-Packs.

- Use-Packs of 1, 5, 10, and 50 sessions
- Use-Packs for unlimited sessions

Note: In SNA Server Version 3.1, the customer had to pay for a maximum of 128 sessions. Sessions above that limit were free. For Communications Server Version 4, a separate, unlimited Use-Pack is provided.

For later increases in your session license requirements, you can use a session-only order. This is a paper-only order—no product shipment will occur.

2.4.2 SNA Upgrade Mechanism

In order to get a fair price for those customers who have already installed any type of SNA software, several upgrade paths are provided. Details should be read in the announcement letter for Communications Server Version 4.

Customers who had SNA Server/6000 Version 2.1, SNA Server/6000 Version 2.2, or any Release of SNA Gateway installed will get an authorization for a number of sessions with their upgrade, depending on the processor group or other functions.

Only sessions ordered during the upgrade are considered as upgrade sessions. Later session upgrades must be done using the normal ordering procedure.

2.4.3 Ordering Sample

Let's assume you have a configuration installed like the one shown in Figure 8 on page 103. The machine TESTSERV (Model 250) has two local HCON sessions, two Gateway HCON sessions from the machine TESTCLI and, in addition, four LU 6.2 sessions are used between the two systems. In total, there are eight sessions required for TESTSERV and six sessions for TESTCLI.

Since both machines are migrating from a different SNA level, you need different upgrade features.

Machine TESTSERV Sample: (Upgrading from SNA Server/6000 Version 2.1)

```
Prod/Feature  Qty   Description
5765-652
A6LE          1 (*) Base upgrade D5 from SNA Server/6000 Version 2.1
A6LP          1     + 16 sessions from SNA Gateway/6000 V 2.1

(*) Base upgrade includes 5 sessions on the D5 processor group.
    Thus the customer can use 5 + 16 sessions after the upgrade.
```

Machine TESTCLI Sample: (Upgrading from SNA Services/6000 Version 1.2)

```
Prod/Feature  Qty   Description
5765-652
A6LD          1     Base upgrade from SNA Services/6000 V1.2
A6LW          1(*)  Session upgrade Use-Pack 5
A6LU          1(*)  Session upgrade Use-Pack 1

(*) This upgrade does not include any free sessions.
    Thus the customer can use 6 sessions after the upgrade.
```

2.4.4 iFOR/LS and License Key Management

Communications Server Version 4 uses iFOR/LS keys to manage the product license. The key is automatically installed when using the official product media. However, Communications Server Version 4 can also be obtained from demo CD-ROMs or from other sources which contain try-and-buy keys which will time out after 60 days.

In order to obtain a permanent key, the user must call an IBM Key Center. More details about the key handling and management are given in 3.5.1, "SNA iFOR/LS Key Handling" on page 85.

2.5 Migrating Other SNA Applications

If you migrate SNA, you also have to check for the status of your SNA applications. To check the availability or prerequisites for migrating any product to AIX Version 4, you should look at *A Holistic Approach to AIX V4.1 Migration, Planning Guide*, SG24-4651; however, you should check the latest updates to the information in the planning guide for SNA products by also examining Table 7 on page 68.

The following list explains the category column of the table.

There are eight categories into which all of the IBM Program Products can be placed when migrating from an AIX V3.2 environment to an AIX V4.1 environment. The categories are as follows:

1 The program product is packaged as part of the AIX Version 4.1 Base Operating System.

2 The program product will run As Is from AIX Version 3.2 product media.

3 The program product will run from AIX Version 3.2 product media when installed with additional program update (PTF).

4 A free upgrade to the same product at a later level is required.

5 A chargeable upgrade to a new product number is required.

6 A free upgrade to a new product number is required.

7 The order of a new product is necessary.

8 A refresh of the software needs to be ordered.

Note: In this instance, the word *upgrade* refers to the pricing structure and does not necessarily correspond to the technical migration action.

In addition to the above listed categories for each program product, you will find an indication of its suitability for the RS/6000 Symmetric MultiProcessor (SMP) models G30, J30 and R30.

U The program product is uniprocessor-safe only (it is not supported on an SMP system).

S The program product is SMP-safe (it will run both on a UP or SMP system).

The SMP information also appears under the category heading. If there is no corresponding SMP information in the table, the information was not available when this document was created.

Table 7. OPP/LPP Mapping - Communications and Drivers

Program Number	AIX V3.2 Program Product	Program Number	AIX V4.1 Program Product	Category	License Key Required
OPP	Block Multiplexer (BLKMUX) Channel Connectivity Option	5697-037	IBM Block Multiplexer Channel Adapter For AIX V1 (WDFM)	5,U	-
		5765-604	Block Multiplexer Channel Connectivity for AIX V1.1	5,S	-
Note: The Block Multiplexer Channel Adapter LPP (5697-037) supports only TCP/IP protocols on the BLKMUX adapter. SNA Server/6000 Version 2.1.2 provided a channel support feature where TCP/IP and SNA traffic are supported. Block Multiplexer Channel Connectivity for AIX V1.1, (5765-604) supports both TCP/IP and SNA protocols and includes SMP support. Customers can upgrade from Block Multiplexer Channel Adapter to Block Multiplexer Channel Connectivity. The Block Multiplexer Channel Adapter LPP (5697-037) will be withdrawn from marketing, effective from December 12, 1995.					
OPP	ESCON Connectivity Option	5765-603	ESCON Channel Connectivity for AIX V1.1	5,S	-
5648-129	Client Input Output/Sockets V2	5648-129	Support available with APAR PN74395	3	-
5765-449	MERVA for AIX V1	5765-449	MERVA for AIX V1	2	-
5696-943	SNA Application Access for AIX V1.1	-	Support planned for 2Q96	-	-
5696-944	SNA Client Access for AIX V1.1	5696-944	SNA Client Access for AIX V1.1	2	-
5696-944	SNA Client Access for AIX V1.2	5696-944	SNA Client Access for AIX V1.2	8,S	-
Note: Client Access V1.2 will be available in February 1996 for AIX 4.1					
5765-233	AIX SNA Manager/6000 V1.1	5765-233	AIX SNA Manager/6000 V1.3	8,S	-
Note: Fix IX52339 (PTF U439966) is required for AIX 4.1.3					
5622-242	NetView FTP Client V1.1 for AIX	5622-242	Support available with PTF U440621	3	-
5765-435	NetView FTP Server V1.1 for AIX	5765-435	Support available with PTF U440619	3	-
5765-247	SNA Server for AIX Version 2	5765-652	Communications Server Version 4 for AIX	5,S	Yes

Chapter 3. SNA Migration Methodology

This chapter covers the migration of AIX systems with SNA communication software installed. It will help you to find the appropriate migration path and to prepare the migration. It will give you all necessary information to achieve a working SNA environment after the migration.

3.1 Migration Path

The migration path for SNA, or for any application using SNA communications, will be shown in this chapter. When migrating SNA-dependent applications, the general migration path below can be used as a guide. Migrating SNA itself depends on the SNA version you are migrating from and will be shown in more detail.

3.1.1 General Migration Path

Migrating any AIX application can be shown in simple steps. This flow can be used for all SNA-based applications, and it is summarized below.

1. Save configuration files
   - Most configuration files are in the /etc directory
2. Save any application installation files
   - These could be installation images in /usr/sys/inst.images
3. Save user data or application-dependent data
   - Including any shell scripts or programs
4. Upgrade the AIX system. This includes upgrading SNA.
   - Use any installation method
5. Customize the AIX system
   - Users, passwords, network, and so on
   - Migrate SNA profiles and re-establish communications
6. Install the application
   - Either a new version, or re-install the old version
7. Customize the application
   - Restore old configuration files
   - Change or migrate the definitions
8. Restore previous user or application data
   - Check for possible changes
9. Test the application

3.1.2 SNA Migration Path

We have tested two possible cases of migrating an AIX Version 3.2 system with SNA to AIX Version 4.1. These cases depend on whether the system has SNA Services/6000 Version 1.2.1 or SNA Server/6000 Version 2.1 installed. The two migration paths are shown in Figure 4 on page 71.

SNA Server/6000 Version 2.1 and Communications Server Version 4 are virtually identical and use the same format for profiles and logs. This means that there are no modifications to make to your SNA configuration during a migration from SNA Server/6000 Version 2.1. SNA Services/6000 V1.2.1 is a totally different version of SNA, with different functions and different file formats used for storing the SNA profiles. This introduces an additional step into the migration process, where we need to update these profiles. This is mostly performed by using a tool provided as part of the Communications Server Version 4 distribution.

We did not perform any testing with levels of SNA older than SNA Services/6000 Version 1.2.1 (sna.sna.obj 1.2.101.315); however, we expect that most of our conclusions will apply equally to earlier versions of SNA Services/6000.

3.1.3 Migrating AIX with SNA Server/6000 Version 2.1 Installed

Although Communications Server Version 4 has some enhancements and changes, from the SNA perspective, and in particular for migration, it is identical to SNA Server/6000 Version 2.1.2.

In this case, SNA does not add many complications to the AIX migration process. You can use the same SNA profiles and API programs used on AIX Version 3.2. The process for migration of AIX with SNA follows a similar path to that of a new Communications Server Version 4 installation.

The recommended procedure is as follows:

1. Back up your SNA profiles and user programs.
2. Install AIX Version 4.1 and Communications Server Version 4.
3. Import the SNA profiles and perform session testing.
4. Set up the SNA API, and test any application programs.

We suggest that for easier problem determination, you should perform testing at each step in the process. These steps are explored in more detail in the following sections.

3.1.4 Migrating AIX with SNA Services/6000 Installed

SNA Services/6000 is the older of the two SNA versions that are supported under AIX V3.2. When migrating a system with SNA Services/6000 Version 1.2.1 installed, you will have to migrate SNA itself. This is a more complicated procedure than simply migrating the AIX operating system and installing a new version of SNA. SNA Server/6000 Version 2.1, SNA Server Version 3.1 and Communications Server Version 4 have major changes from SNA Services/6000. You will need to migrate and then check your SNA profiles, modify your API programs and learn new SNA commands. The recommended procedure in this case is:

1. Back up your SNA Services/6000 profiles, LU 0 profiles and user programs.
2. Install AIX Version 4.1 and Communications Server Version 4.

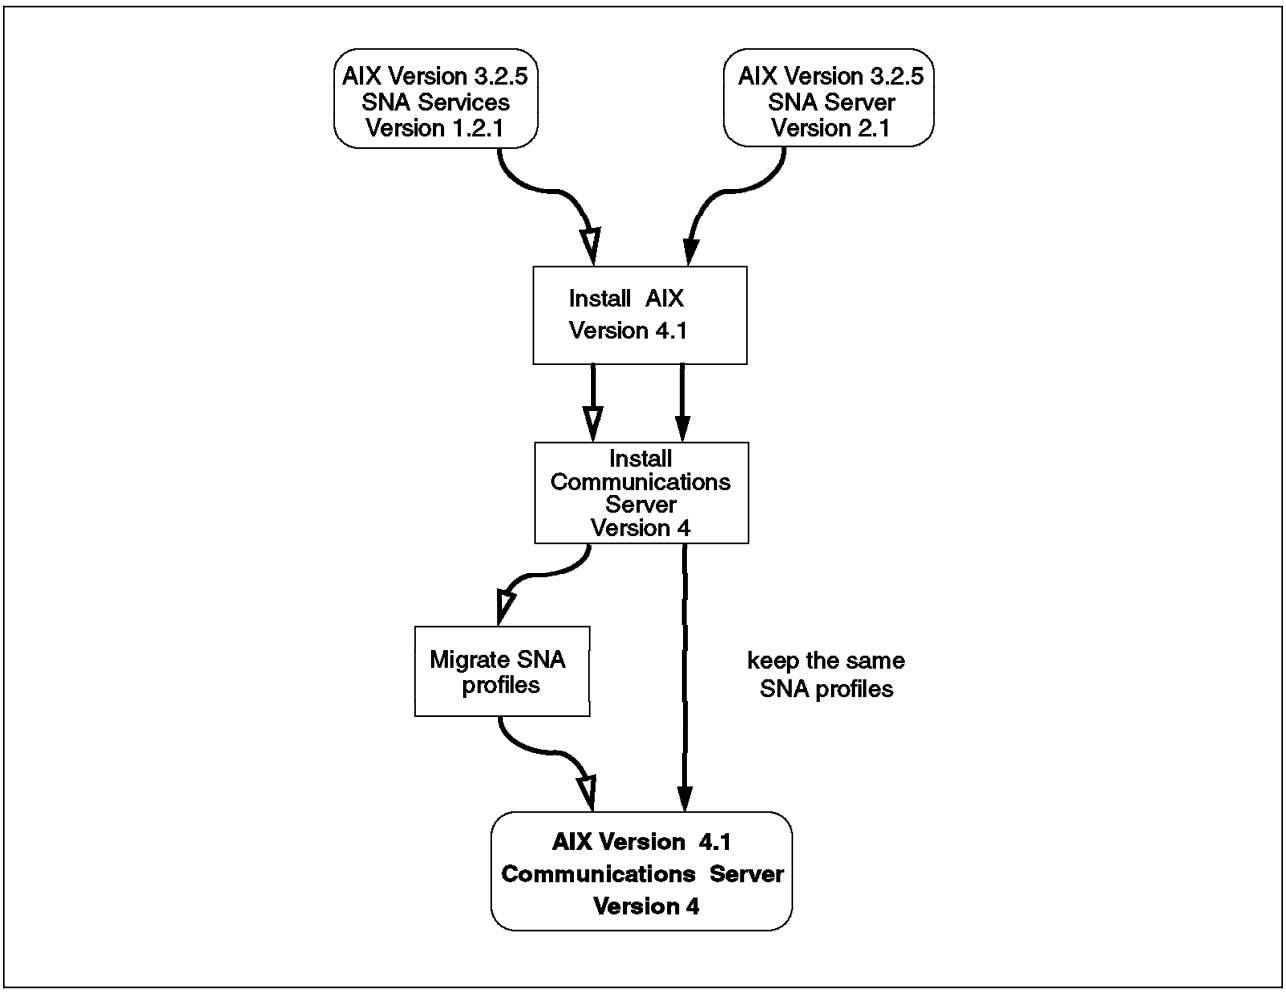

Figure 4. SNA Migration Paths

3. Migrate the SNA Services/6000 profiles and LU 0 profiles to Communications Server Version 4.
4. Import the SNA profiles and perform session testing.
5. Migrate any SNA API programs, if necessary.
6. Set up the SNA API, and test applications.

Again, these steps are explained in more detail in the following sections.

There is a second possible path for performing this migration, as shown in Figure 5 on page 72. In this case, SNA is upgraded to Version 2.1, and the SNA profiles are migrated and tested while still on AIX Version 3.2. When SNA is again operating successfully, then AIX is migrated to Version 4. This method has the advantage of providing a smaller granularity of problem determination, making it easier to isolate the cause of any problems encountered to either the SNA or AIX migration. However, this would require the customer to obtain both SNA Server/6000 Version 2.1 and Communications Server Version 4 and would probably require longer system downtime. In most cases, it will be simpler to perform the migration in one step.

Migration of your SNA Services/6000 profiles is considerable work. The migration command `migratesna` is provided with Communications Server Version 4, but it is not perfect. For example, some profiles no longer exist in the new

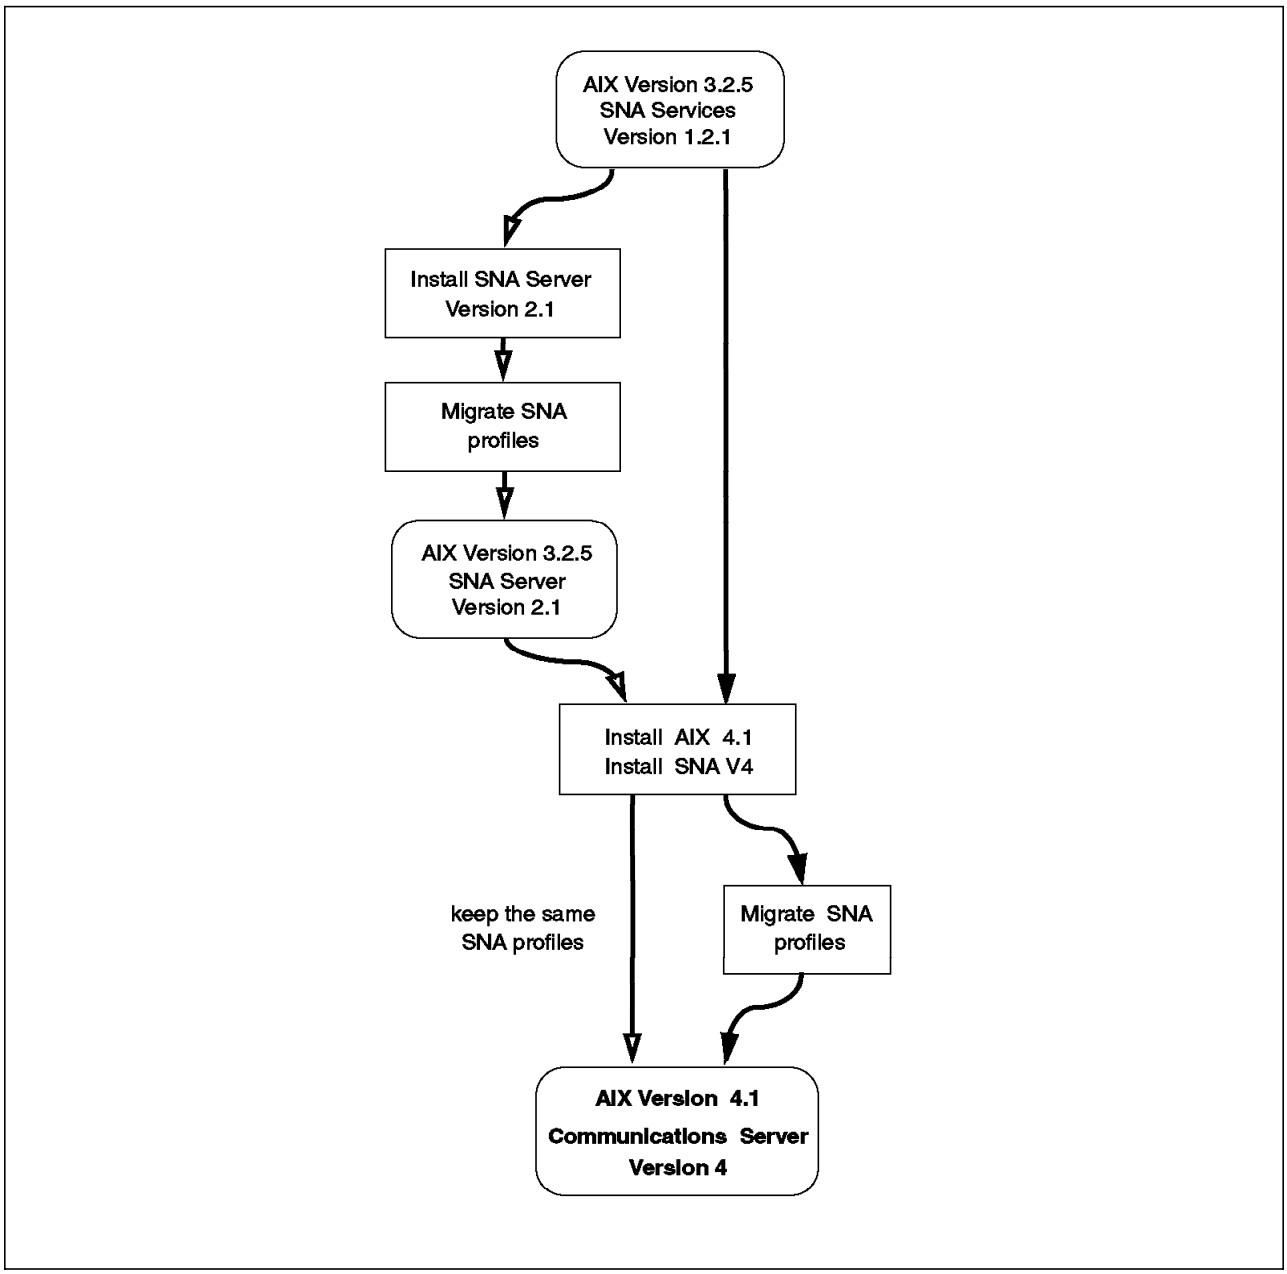

Figure 5. Migrating from SNA Services/6000 Version 1.2.1

version, and many data fields in other profiles have been changed. SNA Server/6000 and Communications Server Version 4 support only a single Control Point (CP). This is in contrast to SNA Services/6000, which allowed multiple CP profiles to be defined even though they were not used in parallel.

The structure and naming of some profiles also changes between versions. In SNA Services/6000 Version 1.2.1, profiles were defined for attachments and connections. In Communications Server Version 4, the terms used are Link Stations and sessions. This is more consistent with the SNA products on other platforms, such as OS/2 and MVS.

When the profiles are migrated, you may find that your profile naming convention is no longer very logical. For example, you may have had an attachment profile called ATT01. You now will find that you have a Link Station profile and an SNA

DLC profile, both called ATT01. You may wish to change the names to be more logical; however, you must be careful to also change any Transaction Programs that depend upon these profiles.

Some programs, like HCON, may point to profile names which no longer exist after the migration. On SNA Services/6000 Version 1.2.1 systems, HCON points to connection profiles, while on SNA Server/6000 Version 2.1 or Communications Server Version 4 systems, HCON must point to LU session profiles. There may be also LU 6.2 applications which had pointed to connection profiles, but now need to point instead to a side information profile. If you know this, it is very easy to name the new profiles like the old connection profiles, and the migration will work immediately. You will not experience these types of problems if you used the *Quick Configuration* method on your old system because the naming convention used by this method was very consistent and, for example, the new LU session profile after migration has the same name as the old connection profile.

A full description of the profile changes is beyond the scope of this document. Complete details of the differences between SNA Services/6000 and Communications Server Version 4 profiles can be found in these documents:

- *AIX SNA Server for AIX User's Guide*, SC31-8211.
- ASKQ DOC ID: G010582 SNA Server/6000 migration guide.
- *IBM AIX SNA Server/6000 Version 3, Release 1 Notes*, found in the file /usr/lpp/sna/bin/README.

Note: For Communications Server Version 4, the titles on most of the documentation still use the name SNA Server for AIX.

3.1.5 Migrating SNA API Programs

Migrating programs written to the SNA Services/6000 or SNA Server/6000 APIs should not be a major problem. Most often, any problems encountered will be due to differences between AIX Version 3.2 and Version 4.1, rather than to changes in the API definitions. For example, the iconv subroutine is commonly used in SNA Transaction Programs (TPs) for data conversion. The value of the return codes sent to indicate successful or unsuccessful completion of iconv changed between AIX Version 3.2 and Version 4.1 as a result of POSIX compliance. Code that uses the iconv routine will probably need to be changed to support this difference.

Binary compatibility with well-behaved SNA Services/6000 Transaction Programs is an objective of SNA Server/6000. But, for some cases of Transaction Programs that use LU 6.2, there are known problems where binary compatibility is not achieved. The details are explained in 3.7, “Migrating SNA Applications” on page 94. Additional information can be found in the *AIX SNA Server for AIX Transaction Program Reference*, SC31-8212.

3.2 Preinstallation Tasks

As usual, the most important task to perform before migrating your system is to back up all critical configuration files. The procedures for backing up AIX system files and for creating a bootable system backup are given in *A Holistic Approach to AIX V4.1 Migration, Volume 1*, SG24-4652. The procedures for backing up SNA- and HCON-specific information are given below. It is also necessary to

choose the type of installation method to be used; however, that decision may be made based upon the best method for the complete AIX installation rather than made specifically for SNA.

3.2.1 Choosing an Installation Method

There are three methods of installation available for AIX Version 4.1. Any of these methods can be used to update a system with SNA installed. The best method depends on your individual installation.

Complete Overwrite Install: In an Overwrite install, all rootvg file systems will be deleted and recreated from scratch. You must be careful to manually back up your data, programs and profiles. You will then need to manually restore or recreate your configuration, and manually restore your data.

The benefit of this approach is that the space used in the final system installation is less because you are able to choose exactly the required filesets. You also have the option of utilizing new features such as modified disk fragment sizes, disk striping and disk compression for user file systems. If your SNA system is also used for server functions, the benefits of these technologies could be considerable.

Preservation Install: This method will remove the contents of /usr, / (root), /var, and /tmp. This includes any non-AIX-system configuration file in the /etc directory. If your programs and data are in /home or any other user-created filesystem, they will be preserved.

During this installation, AIX system files and programs are re-installed; but be aware that all non-AIX-system products installed in /usr/lpp will be deleted, and they will not be re-installed. For example, SNA products like SNA Server and SNA Client Access will no longer exist after the Preservation installation. SNA profiles in /etc/objrepos/sna/ will be lost, and the HCON user definition file /etc/hcon/users will be erased, as will any other configuration data from non-AIX-system applications. Therefore, you must save your profiles to /home or another filesystem or to an external device.

If your system configuration is complicated but you still wish to refresh your AIX file systems and reduce the size of the final system, this method could be suitable.

Migration Install: Migration installation is the easiest and most popular way to migrate an SNA system. This method will preserve your SNA and HCON profiles.

If you are using SNA Server/6000 Version 2.1 on AIX V3.2.5, you need only update the Object Data Manager (ODM) database by using the `verifysna` command after SNA installation. Be aware that some filesets will be installed automatically, based on the Optional Program Products (OPPs) that were installed in your AIX V3.2.5 system.

If you are using SNA Services/6000 Version 1.2.1 on AIX V3.2.5, this method has less advantages. You still have to migrate your SNA profiles manually and check your Transaction Programs. This process will take the same time as in any other installation method. Using the migration path has the advantage of automatically saving all your data but the disadvantage of wasting some disk space.

Which way is best? If you have limited disk space, you are migrating to another system, or want to refresh your filesystem, a Complete Overwrite is recommended. If you want the easiest migration process, a Migration install is probably the best option.

3.2.2 Backing Up SNA Configuration Profiles

Before starting your migration, you should take a backup copy of your SNA configuration profiles. Use the following procedure to export and back up your current SNA profiles:

1. To export the files, enter the following commands on the AIX command line:
   - SNA Services/6000 V1.2:

     `exportsna -f Filename`

     or use the fastpath

     `smit _snaexport`
   - SNA Server/6000 V2.1:

     `exportsna -A -U -f Filename`

     or use the fastpath

     `smit _snaexport`
2. To back up the files, enter the following command on the AIX command line:
   - SNA Service/6000 V1.2 or SNA Server/6000 V2.1:

     `backup -ivf /dev/xxxx Filename`

     Where xxxx is the name of your backup device (for example, rmt0 or rfd0), and `Filename` is the name of the file used in the export commands above.

For SNA Server/6000, the `exportsna` command backs up all configuration profiles, including LU 0 profiles. In SNA Services/6000, however, LU 0 profiles were maintained separately. To back up LU 0 profiles for SNA Services/6000, use the following procedure:

1. Type `lu0config` on the AIX command line, and press **Enter**. You must have root authority to perform this function. You will see the screen shown in below:

```
-------------------------------  MAIN MENU  -------------------------------




                    1. Define LU0 Secondary.
                    2. LU0 Primary.
                    3. Print Configuration File.
                    4. Exit.

                    Enter Option Number :



F3=exit
```

2. Select Option 3, **Print Configuration Profiles**, from the LU 0 configuration menu, and press **Enter**.

   SNA Services/6000 copies all LU 0 profiles to the file /var/lu0/lu0config.rpt.

3. Copy the file /var/lu0/lu0config.rpt to another diskette:

   ```
   cd /var/lu0 ; backup -ivf /dev/xxxx lu0config.rpt
   ```

Attention!

Of the three AIX installation methods, only the Migration installation method will preserve your SNA profiles. It is important that you have your own backup copy.

3.2.3 Back Up HCON Profiles

If you also use 3270 Host Connection Program (HCON) with SNA Services/6000 Version 1.2.1 or SNA Server/6000 Version 2.1 under AIX V3.2.5, you will have to upgrade to HCON Version 2.1 under AIX 4.1.

Backing up the HCON profiles can be a difficult task. The profiles are defined on a per-user basis and are stored in the user's home directory. The user may also have individually tailored their colors and keyboard mappings. The system default values for key tables and color tables are stored in the files /usr/lib/hcon/e789_ktbl and /usr/lib/hcon/e789_ctbl; however, each user can specify their own keyboard and color tables for each HCON session they use. The names of these tables are not restricted to a particular naming convention, and they could be called, for example, $HOME/color_type01 or /etc/hcon_c_map. They may even be in a different directory for sharing with other users. If you miss these files, the HCON session cannot be started and will give the error:

```
0789-020 e789: Cannot open file
```

You must back up these files correctly. This sample script will help you to find the names of the user's color and keyboard table files.

```
#! /usr/bin/ksh
#
# Sample shell script to list HCON keyboard and color table files
# for all HCON users.
#
ODMDIR=/etc/hcon odmget users | grep user_name | tr "\"" " " |
{
  while read junk junk name
  do
      eval export ODMDIR=~$name
      odmget usrdtlts | egrep "ktbl_nam | ctbl_nam"\
          | tr -d "\"" | cut -d"      " -f3
      eval export ODMDIR=~$name
      odmget usrdtlts | egrep "ktbl_nam | ctbl_nam"\
          | tr -d "\"" | cut -d"      " -f3
  done
} | sort -u
```

You should at least back up the following systemwide files:

/etc/hcon/users	Authorized users in the ODM file
/usr/lib/hcon/e789_ktbl	Default keyboard table
/usr/lib/hcon/e789_ctbl	Default color table

You should also back up the following files for each HCON user:

$HOME/usrprofs	Session profiles for the user
$HOME/usrprofs.vc	Session profiles object class
$HOME/usrdflts	Defaults profiles for the user
$HOME/usrdflts.vc	Defaults profiles object class
$HOME/SYS*	HCON autolog profiles

Depending on the type of installation used for AIX Version 4.1 and HCON Version 2.1, these files may be maintained without the need to manually restore them. However, it is always safer to have your own independent copy in case of problems.

3.3 Prerequisites

This section describes the requirements for installing Communications Server Version 4, including disk usage and specific AIX, SNA and HCON fileset requirements.

3.3.1 Disk Space Requirements for Communications Server Version 4

Communications Server Version 4 requires 26 megabytes of free space for installation. The permanent requirement is 19 megabytes; seven megabytes are required only during installation time and are not used after installation. In addition, 0.2 megabytes per language are required for messages.

Table 8 on page 78 shows the itemized disk-space requirements for Communications Server Version 4 components.

Note: Communications Server Version 4 will be shipped with all of these filesets on the media. It is up to the user to install only those parts that are really required.

Table 8. Disk Space Requirements for Communications Server Components

LPP	Description	Space Req.
sna.rte	Communications Server Base	19 MB
sna.lu0	Logical Unit 0 (LU0)	1.2 MB
sna.msg.Lang	Communications Base Messages	0.2 MB
sna.gw	SNA Gateway	3.0 MB
sna.anynet.base	AnyNet Base	0.5 MB
sna.anynet.socksna	AnyNet Sockets over SNA	7.0 MB
sna.anynet.snaip	AnyNet APPC over TCP/IP	1.0 MB
sna.msg.en_US.anynet.rte	Anyet Messages - U.S. English	0.1 MB
sna.toolkit.rte	APPC Application Suite Applications	0.5 MB
sna.toolkit.basic	Basic APPC Connectivity Programs	0.2 MB
sna.toolkit.aftp	APPC File Transfer Applications	0.7 MB
sna.toolkit.aname	APPC/APPN Name Server	0.2 MB
sna.toolkit.3270	APPC 3270 Emulator	0.4 MB
sna.toolkit.misc	Miscellaneous APPC Applications	0.2 MB
sna.snapi	APPC Interactive Application Development Toolkit	0.8 MB
sna.msg.en_US.snapi	SNAPI Messages - U.S. English	0.1 MB
sna.dlcchannel	SNA Channel Support for AIX	3.0 MB
sna.xsna	XSNA Graphical User Interface	1.0 MB
sna.man.en_US.data.rte	SNA Base Manual Pages - en_US	0.3 MB
sna.man.en_US.data.lu0	SNA LU 0 Manual Pages - en_US	0.3 MB
sna.instdlc.ethernet	SNA AIX Ethernet DLC Inclusion Fileset	0 MB
sna.instdlc.token	SNA AIX Token Ring DLC Inclusion Fileset	0 MB
sna.instdlc.fddi	SNA AIX FDDI DLC Inclusion Fileset	0 MB
sna.instdlc.sdlc	SNA AIX SDLC DLC Inclusion Fileset	0 MB
sna.instdlc.x25	SNA AIX X.25 DLC Inclusion Fileset	0 MB
sna.instdlc.channel	SNA AIX Channel DLC Inclusion Fileset	0 MB

Note: The instdlc filesets above require negligable disk space as they do not actually install any significant files. They are simply used through their requisites to ensure that the correct DLC filesets are installed. For more information, see 3.3.4, “Communications Server Version 4 Fileset Requirements” on page 81.

Table 9. Disk Space Requirements for Communications Server Softcopy Manuals		
dtext.brwsr	DynaText Browser	20 MB
sna.books.updoc	Eagle Communications Server for AIX Up and Running Guide	1.4 MB
sna.books.gendoc	AIX SNA General Information Guide	1.2 MB
sna.books.usdoc	SNA Server User's Guide	2.6 MB
sna.books.crdoc	SNA Server Configuration Ref.	3.1 MB
sna.books.cmdoc	SNA Server Commands Ref.	1.7 MB
sna.books.ppdoc	AIX SNA Planning and Performance Guide	1.6 MB
sna.books.dgdoc	SNA Server Diagnostics and Msgs.	2.7 MB
sna.books.gwdoc	SNA Gateway User's Guide	1.4 MB
sna.books.tpdoc	SNA Server TP Reference	4.3 MB
sna.books.cpicdoc	CPI-C API Programming Guide	1.2 MB
sna.books.socksnadoc	AnyNet Sockets over SNA User's Guide	1.3 MB
sna.books.snaipdoc	AnyNet APPC over TCP/IP User's Guide	1.2 MB
sna.books.adoc	APPC Application Suite for AIX User's Guide	1.0 MB
sna.books.chdoc	SNA Server Channel Connectivity User's Guide	1.2 MB
sna.books.escdoc	AIX ESCON Channel Device Driver User's Guide	1.2 MB
sna.books.bmxdoc	AIX Block Multiplexer Channel Device Driver User's Guide	1.2 MB

3.3.2 Disk Space Requirements for HCON

The *3270 Host Connection Program V2.1 and V1.3.2 for AIX: Guide and Reference*, will only give you an overall disk space requirement of 9.4 MB, which includes 3.2 MB for the InfoExplorer documentation. Table 10 shows the space requirements for the different filesets.

Table 10. Disk Space Requirements for HCON Components		
LPP	**Description**	**Space Req.**
hcon.rte	3270 Host Connection Program Base	4.2 MB
hcon.Dt.apps	HCON Desk Top Application	
hcon.Dt.icons	HCON Desk Top Icon	
hcon.X11	HCON for AIXWindows	
hcon.adt	HCON Application Development Tool	
hcon.prod	HCON Productivity Files	
hcon.terminfo	HCON Terminal Definition	0.8
hcon.msg.Lang	HCON Messages	0.2
hcon.loc	HCON Locale	0.4
hcon.info.Lang	HCON InfoExplorer Documentation	3.2 MB

3.3.3 AIX Fileset Requirements

This sections lists some prerequisite filesets that are required when installing Communications Server Version 4. In AIX, you can select to automatically install prerequisite filesets, either in the SMIT installation menu or by using a bundle installation. In these cases, your prerequisite install images have to be on the same installation media as the software being installed.

Compared to AIX Version 3.2.5, the AIX Version 4.1 operating system is divided into smaller installable filesets. The base operating system installation process is much more selective and installs a much smaller set of software by default. Communications Server Version 4 requires vertain filesets beyond those that are automatically installed:

- bos.sysmgt.trace
- bos.rte.ifor_LS

These filesets will not be installed as part of the AIX 4.1 for Clients bundle; however, the bos.rte.ifor_LS fileset will be installed in Server installations. They will both be installed automatically when using Migration installation. Otherwise, they must be installed manually. Communications Server Version 4 will not install if these prerequisites are missing.

The bos.dlc.* filesets, providing the Data Link Control (DLC) interfaces required by SNA, are not installed by default during an AIX upgrade. These files are shipped both on the AIX installation media and on the media with the SNA software. If you choose the easy (or bundle) installation for SNA, or select the sna.instdlc filesets for installation, the appropriate DLC filesets will be installed automatically. Otherwise, you must manually install the DLC for every adapter that will be used by Communications Server Version 4. You can check for the required DLC files with the command:

```
lsdev -Cc dlc
```

For example, if you use SNA on a token-ring adapter, and the DLC fileset is installed, you will see the following message:

```
dlctoken Available Token-Ring Data Link Control
```

If the DLC is not present, it can be installed manually by selecting the following filesets:

- bos.dlc.com
- bos.dlc.*adapter_type*

For more information on DLC fileset installation, see 3.3.4, “Communications Server Version 4 Fileset Requirements” on page 81.

Many customers use the iconv subroutine with Communications Server Version 4 APIs to convert data between different codesets. If you want to convert data on an AIX system, you should install these filesets, which will not be installed by default:

- bos.iconv.com
- bos.rte.iconv
- bos.iconv.*locale*

Where *locale* is the locale for which conversions will be made. For example, bos.iconv.fr_FR.

3.3.4 Communications Server Version 4 Fileset Requirements

As shown in Table 8 on page 78, Communications Server Version 4 is split into many separate filesets. The installation can be performed automatically by selecting the SMIT **Easy Install** path (bundle installation). Alternatively, you can use **Custom Installation** to allow a specific subset of the filesets to be selected.

If the Easy Install path is used, the following filesets will be installed automatically.

- sna.rte
- sna.lu0
- sna.gw
- sna.anynet.base
- sna.anynet.socksna
- sna.anynet.snaip
- sna.toolkit.rte
- sna.toolkit.basic
- sna.toolkit.aftp
- sna.toolkit.aname
- sna.toolkit.3270
- sna.toolkit.misc
- sna.snapi
- sna.xsna
- sna.man.en_US.data.rte
- sna.man.en_US.data.lu0
- sna.instdlc.ethernet
- sna.instdlc.token
- sna.instdlc.fddi
- sna.instdlc.sdlc
- sna.instdlc.x25
- sna.instdlc.channel
- hcon.rte
- hcon.X11
- hcon.Dt.apps
- hcon.Dt.icons
- hcon.prod
- hcon.adt
- hcon.util
- hcon.terminfo.ibm.data
- hcon.terminfo.bull.data
- hcon.terminfo.wyse.data
- hcon.terminfo.dec.data

The corresponding message filesets will be installed automatically by the installation programs.

The sna.instdlc filesets are a mechanism used for automatically installing those DLC filesets for which corresponding device drivers are found on your system. Each sna.instdlc fileset has installed-requisites of the device driver filesets for each adapter of that type. An installed-requisite indicates that the fileset should be installed automatically if the nominated installed-requistite fileset is already installed or is on the list of filesets to be installed. The fileset then has a

co-requesite of the corresponding DLC fileset. A co-requesite is a fileset that must be installed for the fileset to function successfully. For example, the fileset sna.instdlc.token has installed-requesites for:

- instreq sna.rte

and for one or more of the following:

- instreq devices.mca.8fa2.rte
- instreq devices.mca.8fc8.rte
- instreq devices.isa.PNP80CC.rte
- instreq devices.pci.14101800.rte
- instreq devices.pcmcia.a4001e00.rte

This means that if the sna.rte code and one or more token ring device drivers are installed, then the sna.instdlc.token fileset will be installed. This fileset also has the following co-requesite:

- coreq bos.dlc.token

This co-requesite specifies that the DLC fileset must also be installed.

Note!

If you are short on disk space, you should use the Custom Install method to select only those filesets you really need.

The minimal SNA installation could be as small as:

- sna.rte
- sna.msg.*language*
- bos.dlc.*adapter_name*

However, this will depend greatly upon the functions you require.

3.3.5 HCON Fileset Requirements

You may use the 3270 Host Connection Program (HCON) to provide 3270 terminal emulation sessions to your host systems. HCON V2.1 is required on AIX V4.1; so you must also upgrade your HCON level.

The HCON product is also split into separate filesets, but overall, HCON does not use so much space that you need to take special attention on its installation. Only when you are extremely limited on disk space, should you select single filesets.

A single-user authorization for the HCON software is now bundled with Communications Server Version 4. If you select the bundle installation for SNA, this HCON support will be installed automatically. If you wish to extend this support to a full multiuser version, you must place an order for the IBM 3270 Host Connections Program (5765-398).

3.4 Installing the New Software

The installation of the AIX base operating system has been well covered in *A Holistic Approach to AIX V4.1 Migration, Volume 1*, SG24-4652; so in this section, we will consider only the installation of Communications Server Version 4 and HCON.

3.4.1 Installing Communications Server Version 4

As discussed above, when installing Communications Server Version 4, you have the choice of two methods.

3.4.1.1 Easy Install

To install the Communications Server Version 4 bundle you can follow the SMIT Easy Install path, and select the **Media-Defined** bundle.

```
                          Install Bundles of Software

Type or select a value for the entry field.
Press Enter AFTER making all desired changes.

                                                        [Entry Fields]
* INPUT device / directory for software              [/usr/sys/inst.images]   +


  +--------------------------------------------------------------------------+
  |                           BUNDLE to install from                         |
  |                                                                          |
  | Move cursor to desired item and press Enter.                             |
  |                                                                          |
  |   App-Dev                                                                |
  |   Client                                                                 |
  |   DCE-Client                                                             |
  |   Media-Defined                                                          |
  |   Pers-Prod                                                              |
  |   Server                                                                 |
  |                                                                          |
  | F1=Help                 F2=Refresh              F3=Cancel                |
F1| F8=Image                F10=Exit                Enter=Do                 |
F5| /=Find                  n=Find Next                                      |
F9+--------------------------------------------------------------------------+
```

This will automatically install the fileset Communications.Bnd, which is an auto-install bundle. Once this bundle is installed, the software that it specifies will then be installed automatically. The contents of the Communications.Bnd file are listed in 3.3.4, “Communications Server Version 4 Fileset Requirements” on page 81 above.

3.4.1.2 Custom Install

You can also choose to install single filesets by following the SMIT Custom Install path.

Note: Remember that although you are installing Communications Server Version 4, the filesets will show in the installation screens as level 3.1.1.

```
                Install New Software Products at Latest Level

Ty+-----------------------------------------------------------------------------+
Pr|                           SOFTWARE to install                                |
  |                                                                              |
  | Move cursor to desired item and press F7. Use arrow keys to scroll.          |
* |     ONE OR MORE items can be selected.                                       |
* | Press Enter AFTER making all selections.                                    |+
  |                                                                             |+
  | [MORE...24]                                                                 |+
  |     + 3.1.1.0  Miscellaneous APPC Applications                              |+
  | >   + 3.1.1.0  SNA Gateway                                                  |
  | >   + 3.1.1.0  SNA Server Base (LU1 LU2 LU3 LU6.2)                          |+
  |     + 3.1.1.0  SNA Server SNAPI TP development tool                         |+
  | >   + 3.1.1.0  X-windows Management Tool for SNA Server                     |+
  |                                                                             |+
  | > 3.1.1.0  sna.books                                                   ALL  |+
  |     + 3.1.1.0  AIX Block Multiplexer Channel Device Driver Users Guide      |+
  |     + 3.1.1.0  AIX ESCON Channel Device Driver Users Guide                  |
  | [MORE...26]                                                                 |
  |                                                                             |
  | F1=Help                F2=Refresh             F3=Cancel                     |
F1| F7=Select              F8=Image               F10=Exit                      |
F5| Enter=Do               /=Find                 n=Find Next                   |
F9+-----------------------------------------------------------------------------+
```

Caution

Do not select **all_licensed** or **sna ALL** if you are not really sure you want to install all parts of SNA, including the documentation. This will cost you approximately 80 MB of disk space.

In order to install any of the online SNA publications, the DynaText browser must be installed first, or concurrently, by selecting the fileset:

- **2.3.0.2 DynaText Browser**

The SNA books can be selected either as **sna.books ALL** or by selecting a single book, such as:

- **3.1.1.0 AIX SNA General Information Guide**

3.4.2 HCON Installation

If you selected **bundle installation** for Communications Server Version 4, HCON will be installed automatically because it is now bundled with Communications Server Version 4.

You can also install each of the products separately because you might not want to install all of the HCON filesets. For example, the hcon.Dt.rte fileset is required only if you want to use HCON with the Common Desktop Environment.

If you have to add a user or rebuild all your users from scratch, the HCON definitions in the user′s HOME directory will be overwritten by the defaults. In this case, you will now have to restore the definitions from the backup file you made in 3.2.3, “Back Up HCON Profiles” on page 76.

3.5 Post-Installation Tasks

This section details the steps that you must take after installing the software to complete the tailoring of your new environment.

3.5.1 SNA iFOR/LS Key Handling

As written earlier in this book, an iFOR/LS key will be automatically installed during the SNA product installation. However, it is possible that you had a try-and-buy version of the product with a key that will time out after 60 days. You will not get any message when the license expires; however, if you then try to restart your Communication Server, you will see the message:

```
iFOR/LS: License not found in data base
```

and the product will not start again.

In order to use the product again, you must call an IBM Key Center for a permanent key.

In the simplest case, when you obtain the permanent key from IBM, you simply have to edit the /usr/lib/netls/conf/nodelock file. Type in the code exactly as received. Then you can start the SNA product again by using the `sna -start` command.

Refer to *A Holistic Approach to AIX V4.1 Migration, Volume 1*, SG24-4652, or to the iFOR/LS documentation for information on iFOR/LS.

3.5.2 SNA Session Count

When you order the product following the directions in 2.4, “SNA Product Ordering” on page 64, it is necessary to specify a number of sessions. When the product is installed, it is enabled for a single user. It is up to the system administrator to change the number of allowed sessions according to the license purchased.

The root user is allowed to change the *Maximum number of licensed sessions* within the *SNA Node Profile* to the licensed number of sessions.

You can easily check the number of sessions in use and the number of licenses you have entered in the SNA Node Profile by using the command `sna -d sc`, which results in the following output:

```
SNA Session License Count Information

Active end point sessions:              0
Active intermediate sessions:           0
Active secondary lu0 sessions:          0
Active gateway sessions:                0
-----------------------------------------
Total:                                  0

Number of licenses:                     100
Number of sessions exceeding licenses:  0

There have been '0' violations of the maximum
number of licensed sessions since SNA was started.
```

If you exceed the number of sessions, nothing will happen. The sessions will start, and the violation will be shown only by using the above command. The license limit is not enforced by the software. However, you will be breaking your licensing agreement. If you require more licenses, you should contact your IBM representative.

3.5.3 SNA Log, Trace and Dump Sizes

By default, Communications Server Version 4 will be set to use a log file size of 1.5 MB and a file wrap limit of 10. The files are stored in the /var/sna directory. By default, the /var file system is only 4 MB; so after a short time, you will have filled up the /var file system. To avoid this, you should either expand the /var file system or decrease the SNA log file size.

We recommend both expanding the /var file system and decreasing the SNA log file size and wrap limit, unless you have trouble with SNA and are told by software support to change your log and trace file sizes.

To change the log and trace file size and wrap limit, you can use the command:

```
sna -setlogs
```

or use

```
smit _snasetlogs
```

Whenever SNA ends abnormally or is cancelled by the user, an ABEND dump will be written into /var/sna. The size of this SNA_ABEND.dmp file is 1 MB. You should take care to expand the /var file system enough to hold all log and error information, including an ABEND file.

3.5.4 SNA Start Up

In general, SNA will be started automatically by `init`. This calls the script `/etc/rc.sna` where the user must uncomment the real start command by deleting the # comment sign:

```
# Start the Communications Server
#
# To have SNA start at IPL time, uncomment line below
/usr/bin/sna -start
```

This file is also a good place to start any SNA specific resources or applications. Remember that Communications Server Version 4 allows links to automatically start when SNA starts; so a link start here may be obsolete.

3.6 Migrating SNA Profiles

Migration of SNA profiles depends on the installation method used and the SNA version you are upgrading from. The following table will show you the necessary actions for each of the different varieties.

Table 11. SNA Migration Actions

Old Product	Installation Method		
	Overwrite	**Preservation**	**Migration**
SNA Services/6000 V1.2.1	migratesna importsna verifysna	migratesna importsna verifysna	migratesna importsna verifysna
SNA Server/6000 V2	importsna verifysna	importsna verifysna	verifysna

Note:

1. migratesna migrates the old profiles.
2. importsna imports the profiles.
3. verifysna verifies the profiles.

The steps involved in performing these actions are described below. For a detailed description of all SNA migration options and profile changes, please refer to the manual *SNA Server for AIX User's Guide*, SC31-8211.

To understand the migration process, Figure 6 on page 88 shows you the principles of how the SNA profiles are stored and how they can be updated.

3.6.1 SNA Services/6000 Profile Migration Process

Since the structure of the SNA profiles changed totally and many of the profile names changed, the old profiles must be migrated to the new format. Communications Server Version 4 provides a tool to convert the old profiles to the new standard.

If your SNA Services/6000 was older than V1.2.1, you must first convert your profiles to the V1.2.1 format, using the following command:

```
awk -f /usr/lpp/sna/bin/sna_update.awk SourceFile > TargetFile
```

The TargetFile can then be used as the input file for the migration described below.

In some cases, you will not get back an error-free, running system from this migration. Many SNA Services/6000 profiles are retired during the migration; other profiles are allowed only once in SNA Server, while SNA Services/6000 allowed multiple profiles. These changes cannot be expected to be error free without manual intervention. However, if you have a look to the new and old profiles in Figure 7 on page 91, you will have an idea of the profile mapping performed during the migration and will be able to resolve errors.

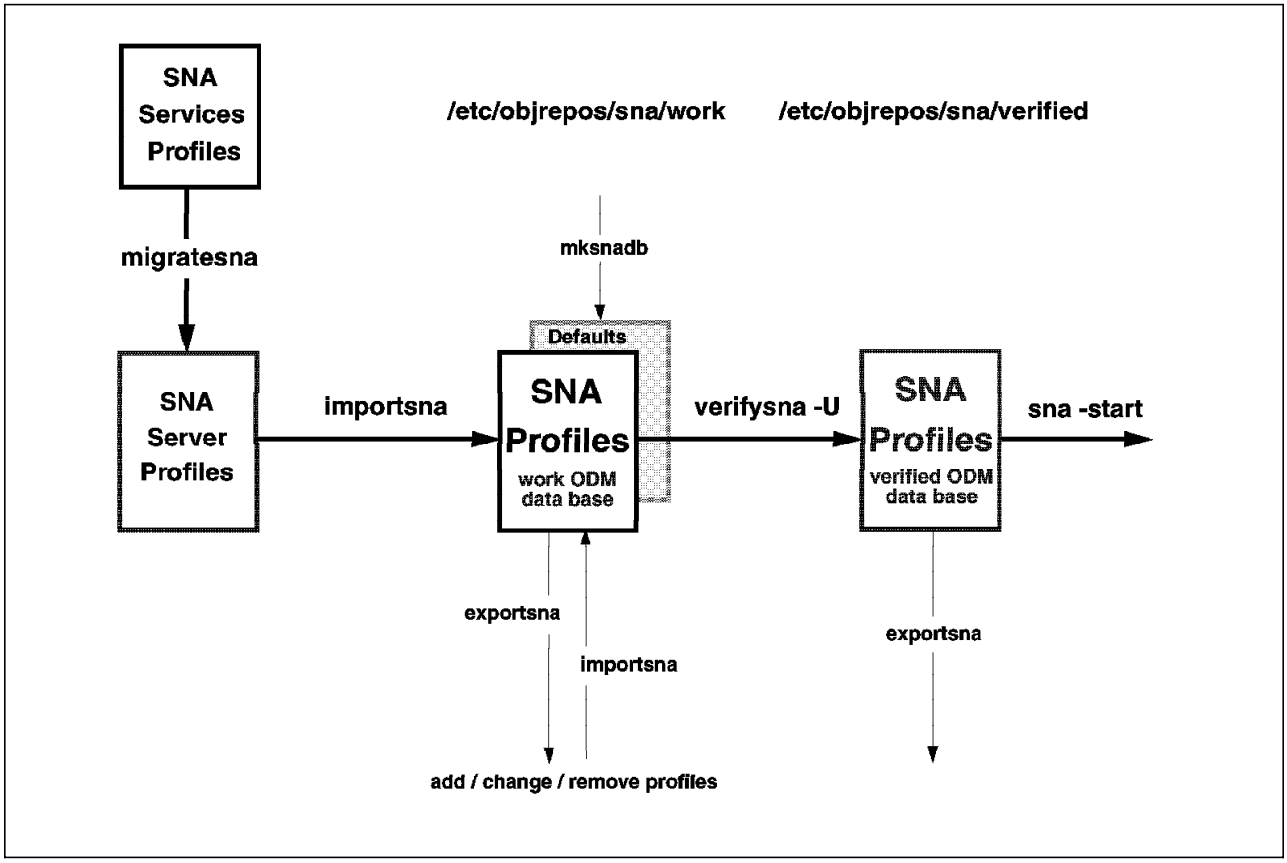

Figure 6. Principle of SNA ODM Data Base

Hint

Edit your old SNA Services/6000 profiles before before running `migratesna`, and remove all unused or *test* profiles. This will help you to get a working set of profiles in a the shortest possible time.

The migration itself should be done in the following sequence:

1. Decide a name for your Control Point (CP). This must be a unique name within your SNA network. It is used for APPN only. If you do not use APPN, any uppercase word up to eight characters in length can be used. Your SNA system programmer may wish to allocate your CP name along with the network name used in the enterprise.
2. Migrate the profiles by using the command:

   ```
   migratesna -s OldProfile -t NewProfile -l NetName.CPName \
       -e ErrorFile
   ```

   or use

   ```
   smit migratesna
   ```

```
                        Migrate Configuration Profiles

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

                                                      [Entry Fields]
* Source file of profiles to migrate                 [OldProfile]
  Target file for migrated profiles                  [NewProfile]
  Error file for invalid profiles                    [ErrorFile]
  Message log file                                   []
  Fully qualified local control point (CP) name      [NetName.CPName]
  Automatically import migrated profiles?             no                    +
  Suppress messages?                                  no                    +
  Replace output files if they already exist?         no                    +
  Verbose mode?                                       no                    +

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

Any profiles that cannot be migrated will be saved in the named ErrorFile. If you later correct the errors, you can re-migrate the all original profiles, or migrate just the profiles that gave errors by using the edited ErrorFile as input to the migratesna command.

3. Repeat the step with your *OldLU0Profiles*, if any. Remember, you may have saved these profiles as lu0config.rpt.

   ```
   migratesna -s OldLU0Profile -t NewLU0Profile -l NetName.CPName \
       -e LU0ErrorFile
   ```

4. Resolve errors and discrepancies, if any.

   Examine the error files for any errors. You can edit the original profiles and re-run migratesna as many times as you like.

5. Import the profiles by using the commands:

   ```
   importsna -f NewProfile
   importsna -f NewLU0Profile
   ```

   Or use smit importsna

```
                        Import Configuration Profiles

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

                                                      [Entry Fields]
* File from which to import                          [NewProfile]
  Replace duplicate profiles?                         yes                    +
  Security information source file                   []

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

6. Verify the profiles by using the command:

   `verifysna -U`

   Or use `smit verifysna`

```
                        Verify Configuration Profiles

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

                                                      [Entry Fields]
  Update action if verification successful            normal_update          +
  If normal_update or dynamic_update,
     Backup file for committed database               []
     Backup security file for committed database      []

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

Keep in mind that even if the profiles migrate with no errors, they still may not work. Some possible reasons for failure are:

- You had multiple Control Points (CP) or Attachment profiles under SNA Services. The `migratesna` program will use just the first CP profile—further profiles will be listed, and you may choose the right one. The Exchange Identifier (XID) used previously in the CP profile will now be within the new Link Station profile. Check the XID; it may be the wrong one, or the field may be blank.

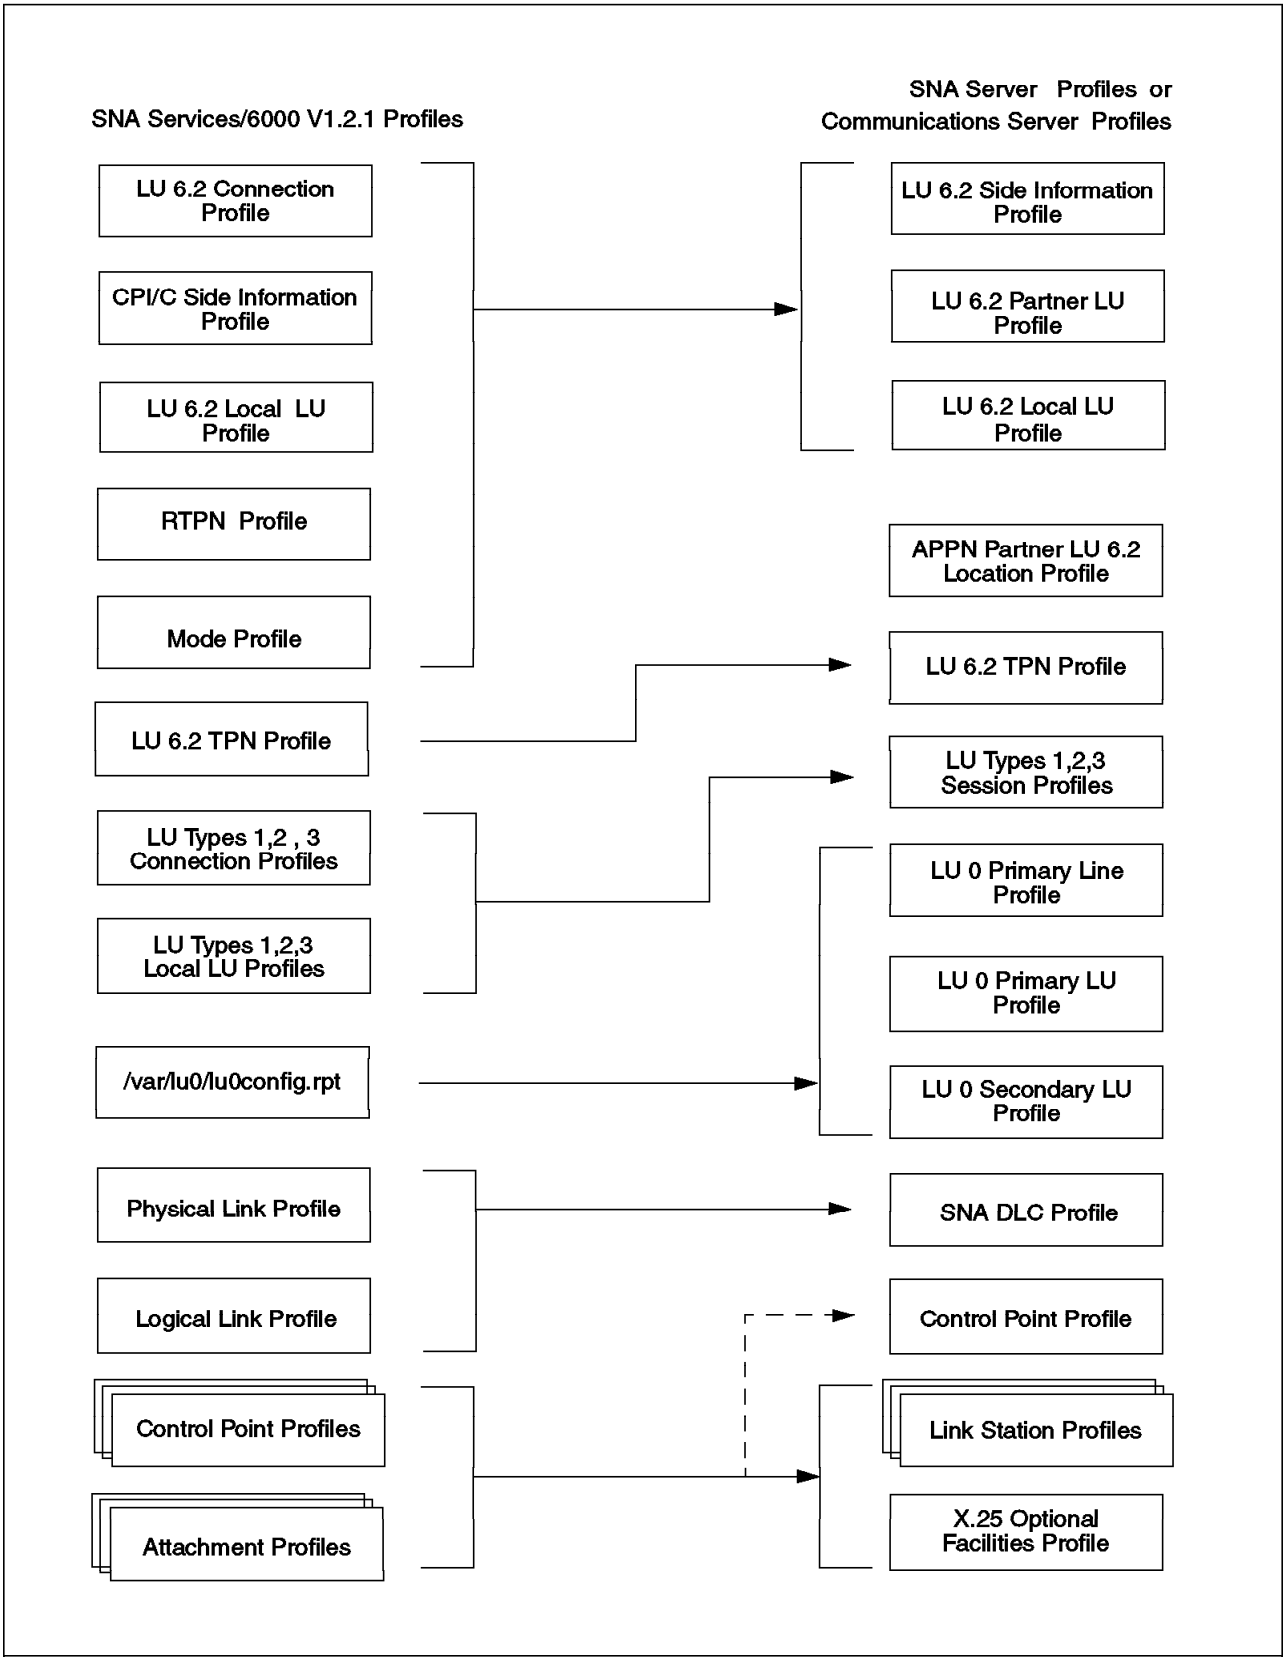

Figure 7. Migration of Information for Communications Server Profiles

- You had multiple LU 6.2 Mode profiles with the same content but different names. The migratesna command will use only the first one—all others will be removed.

In general, most migrations work without problems, but in cases where many errors occur, it might be easier to completely redefine your SNA profiles instead of wasting a long time debugging.

3.6.2 Refresh Default Profiles

To refresh all profiles and start from a clean system, enter the following commands on the AIX command line:

1. Save your definitions for later use.

   ```
   exportsna -A -f Filename -U
   ```

2. Remove the work data base.

   ```
   rm /etc/objrepos/sna/work/*
   ```

3. Remove the verified data base.

   ```
   rm /etc/objrepos/sna/verified/*
   ```

4. Make the default profiles.

   ```
   /usr/lpp/sna/bin/mksnadb
   ```

3.6.3 Import Profiles from SNA Server/6000 Version 2.1

If you had SNA Server/6000 Version 2.1 installed on your old system, you do not need to migrate the SNA profiles. Depending on the installation method, you may only have to import your old definitions and verify them. Your communications should then work as before.

The import process is simple. Restore the backed up files to your disk (if neccesary), and enter the following command on the AIX command line:

```
importsna -f Filename
```

or use

```
smit importsna
```

After that, you must verify and update the SNA database in order to use the new profiles:

```
verifysna -U
```

or use

```
smit verifysna
```

In the case of a Migration installation, your old SNA Server/6000 Version 2.1 profiles will still be in place. The only step needed in this situation is the last verfication step.

3.6.4 Connection Tests

After verifying your profiles, you should then do some testing to check the new system. This process will soon show you if the migration of the SNA profiles was successful or not.

3.6.4.1 Link Station Test

The first step is the Link Station test. A link is the combination of hardware and software that provides direct communication between adjacent nodes. The hardware component of the link consists of an adapter at each node and the transmission medium that connects the two adapters. The software component provides control of the link and the data exchanged over it. The Link Station test allows you to check that the migration process was successful at the link level and that the Initial node, CP, SNA DLC, and Link Station profile parameters are correct. The procedure is as follows:

1. Start the SNA subsystem.

   ```
   sna -start
   ```

2. Start the Link Station.

   ```
   sna -s l -p Link_station_profile &
   ```

3. Check the Link Station state.

   ```
   sna -d l
   ```

After starting the Link Station, you should see a message like:

```
0105-2723 The "austin" Link Station has been started
```

If you do not see a similar message, or you see any error messages, check these profiles:

- Control Point profile (XID, Network name, CP name)
- SNA DLC profile (Device name, Service Access Point (SAP) address)
- Link Station profile (XID, DLC profile name, Remote link address, Remote SAP address, link activation parameters)

Other possible causes are that the DLC device is not available or the remote node might be set wrongly. The effective way to troubleshoot is to start a link trace to see what's going on at the link level. The procedure is as follows:

1. Set the trace link parameter in the Link Station profile to *yes*.
2. Try again to connect the link by using the command:

   ```
   sna -s l -p Link_station_profile &
   ```

3. Stop the link by using:

   ```
   sna -stop l -profile LinkStationProfileName -type forced
   ```

   You will now see the trace file /var/sna/LinkStationProfileName.

4. Format the trace file by using the command:

   ```
   snaformat -A /var/sna/LinkStationProfileName
   ```

   The format command will create a file, LinkStationProfileName.det, which can be viewed by using any editor or by just using the pg command. The output contains all data and control commands from the line being traced. It shows the values as hex values or as formatted SNA terms.

3.6.4.2 Session Test

The next step is the session test. Communication between a Transaction Program and the SNA network occurs through *Network Accessible Units* or NAUs, which are unique network resources that can be accessed (through unique local addresses) by other network resources. NAUs communicate with NAUs in other nodes over temporary logical communication channels called *sessions*. Before two Transaction Programs can communicate, their Logical Units must establish a session.

A successful session test will prove that all of your SNA profiles are correct and are ready for your Transaction Programs. The procedure is as follows:

1. Start the LU 6.2 session:

   ```
   sna -s session -ln local_lu_name -pn partner_lu_name -m Mode_name
   ```

2. Check the LU 6.2 session state by using the command:

   ```
   sna -d s
   ```

3. In the case of LU types 1, 2, and 3, sessions will be started automatically by the remote host after starting the Link Station. You can skip the start session process, and just check the LU 1, 2, and 3 session states with this command:

   ```
   sna -d s123
   ```

The SNA migration program, `migratesna`, could have made some changes within your profiles which cause the sessions to fail.

If there is a problem, check your new profiles again and compare them with the old ones—especially with these profiles:

- Local LU profiles
- Partner LU profiles
- Mode profiles
- Side Information profile
- LU 1, 2 and 3 session profiles

3.7 Migrating SNA Applications

This section explains how to change programs written to use the SNA APIs in SNA Services/6000 to use the updated APIs in SNA Server/6000 Version 2.1, SNA Server Version 3.1 and Communications Server Version 4. It also gives some other notes on SNA Transaction Program migration.

3.7.1 How the APIs Have Changed

Communications Server Version 4 has virtually the same programming interfaces as SNA Server/6000 Version 2.1; so SNA Server/6000 Version 2.1 users will be able to use the same API programs on AIX V4.1. However, Communications Server Version 4 interfaces are different from those in SNA Services/6000 Version 1.2.1. Most of the API changes are provided to support APPN capabilities, such as dynamic link definition, dynamic route computation, topology service, directory services, and intermediate session routing. These API changes enable remote resources, such as Control Points, Logical Units,

Transaction Programs, and associated links, to be supported at run time so that those remote resources do not need to be configured on the local node.

Communications Server Version 4 also supports the Common Programming Interface for Communications (CPI-C) API Version 2.0 conformance classes. A thread-safe version of the CPI-C API is also shipped with Communications Server Version 4 to allow the development of Symetric MultiProcessor (SMP) safe Transaction Programs.

3.7.2 Binary Compatibility

Binary compatibility between AIX Version 3.2 and Version 4.1 extends to well-behaved programs that use SNA Services/6000 and SNA Server/6000 V2.1 published and supported interfaces only. In pursuit of binary compatibility, the differences between the SNA Services/6000 API and SNA Server API are very small:

- Data structure sizes have not been changed.
- Some fields in these structures have been retired but not reassigned.
- Some new structures have been introduced.

3.7.2.1 Exceptions to Binary Compatibility

Binary compatibility with well-behaved SNA Transaction Programs from AIX V3.2.5 is an objective of Communications Server Version 4, except for the following known conditions that affect binary compatibility with SNA Services/6000 LU 6.2 Transaction Programs.

Note: This section applies only to users migrating from SNA Services/6000 on AIX Version 1.2.

SNA Services/6000 Operating Systems Subroutines Limited Interface: The Limited Interface does not exist in Communications Server Version 4. Programs written to the Limited Interface must be re-coded and recompiled to use one of the Application Programming Interfaces that are supported by Communications Server Version 4.

Remote Transaction Program Name (RTPN) Profiles: Some SNA Services/6000 LU6.2 Transaction Programs use the RTPN profile to provide one or more of the following parameters for the `allocate` command:

- Remote transaction program name
- Conversation type
- Sync level

Because the RTPN profile is not present in Communications Server Version 4, these parameters must be defined elsewhere. The first of these, Remote Transaction Program name, can now be defined in the Communications Server Version 4 Side Information profile. However, Communications Server Version 4 profiles do not provide fields for Remote Transaction Program Conversation type and Sync level. To migrate Transaction Programs using these fields, change the source code to specify a valid Conversation type and Sync level, and recompile them.

Session Reconnect: Support for this feature has been removed. If you use it, you must change the source code and recompile your program.

LU 6.2 writex and readx with allocate: This function is no longer supported for LU 6.2. If you use it, you must change the source code and recompile.

Data structure changes: These fields have been retired. If you use them, you must change the source code and recompile your programs.

- ext_io_str
  - allocate - retired for LU 6.2 only
  - priority
  - tpn_option
  - usrhdr_len
  - usr_trunc
- allo_str
  - priority
  - recov_level
- attr_str
  - conn_status
  - recovery_level

Opening Side Information profiles: In SNA Services/6000, programs using the Operating System Interface or Library Subroutines APIs would open a Connection profile to initialize the interface to SNA. With Communications Server Version 4, however, programs initialize access to SNA by opening a Side Information profile. As a feature of Communications Server Version 4, a program can open a Side Information profile that has not been previously defined. The program is then responsible for specifying the parameters in the allo_str structure that would have been specified in the Side Information profile:

- Local LU Name
- Partner LU Name
- Mode Name
- Remote Transaction Program Name

If a program opens an undefined Side Information profile, and does not specify the required parameters in the allo_str structure, the subsequent `allocate` will fail with an error SNA_PARMS (146), which indicates that invalid parameters were passed in the `allocate` call.

3.7.3 Profile Changes from SNA Services

Your Transaction Programs will use profile names in the `init` or `allocate` functions. You should check your profile names after using the `migratesna` command because they may have changed. If your Transaction Programs are hard-coded to use specific profile names, you may need to modify them. Some Transaction Programs use the *Connection* profile to identify the remote LU. They now have to point to the *Remote LU* or *Side Information* profile to identify the remote side.

3.7.4 SNA Command Changes

You may have shell scripts to provide status reporting, to start or stop resources, or to control traces for SNA by using the System Resource Controller (SRC) interface. Although Communications Server Version 4 continues to support the SRC interface (lssrc, startsrc, stopsrc, traceson, and tracesoff), some output formats may have changed.

The recommended interface is the `sna` command which provides all necessary options for controlling SNA resources:

- `sna -start`
- `sna -stop`
- `sna -display`
- `sna -trace`

See the *SNA Server for AIX Command Reference*, SC31-8214, for the complete details of command syntax and usage.

Before migrating a shell script for use with Communications Server Version 4, you should compare the functionality of that script with new features provided as a part of the new version. For example, shell scripts which would initiate recovery procedures in case of link failures are not necessary since Communications Server Version 4 allows automatic restart or recovery under the following conditions:

- SNA start
- Normal deactivation
- Abnormal deactivation
- Link is required by the Control Point (APPN)
- Link establishment requested by a Transaction Program

3.7.5 Recompile with C for AIX (New C Compiler)

You need to recompile your Transaction Programs after modifying them for profile parameter changes, command changes or for other reasons. The XL C compiler used on AIX V3.2 is not supported on AIX V4.1. You should use the C for AIX Licensed Program Product. There should be no problems related to SNA with the new compiler; however, be aware that the C Compiler is no longer bundled with the operating system. You have to order it separately.

3.7.6 General Changes

In AIX V4.1, some libraries and commands have changed. These changes may effect your Transaction Programs. Check the following changes to see if any will effect your programs.

3.7.6.1 iconv

Often, Communications Server Version 4 is used for communication with other products and other platforms, and many SNA platforms use the EBCDIC data format; so you often have to use the `iconv` subroutine in your Transaction Programs.

File Converters: Not all of the `iconv` filesets you need will be installed automatically. You have to check and install file converters manually; otherwise `iconv_open` will return an EINVAL error. See 3.3.3, “AIX Fileset Requirements” on page 80 for details about the necessary filesets.

iconv Return Values: The values returned by `iconv` have changed to conform to XPG4. In its AIX Version 3.2 implementation, the `iconv` subroutine returned the number of bytes remaining in the input buffer on error, or 0 on success. In AIX V4.1, `iconv` returns -1 on errors (and sets errno) and >=0 on successful completion. For `iconv`, a non-zero no longer indicates an error.

AIX V3.2 binaries will continue to run and operate correctly on AIX V4.1. Recompiling AIX V3.2 code with AIX V4.1 should complete successfully. However, due to changes in the values returned, your program may behave differently and may need to be modified.

iconv_close Return Values: The values returned by `iconv_close` have also changed. Now, *iconv_close* returns -1 on errors (and sets errno) or 0 on successful completion. It previously did not return anything. This change was made for conformance to XPG4. It should not affect applications.

3.7.6.2 AIX Command Changes

Some AIX commands have changed in AIX V4.1. Most of the changes were made for standards compliance. In some cases, these changes can effect the behavior of your shell scripts or Transaction Programs. You will find an overview of the changes in *A Holistic Approach to AIX V4.1 Migration, Volume 1*, SG24-4652.

3.7.7 Application Test

Normally, SNA Transaction Programs will behave well on AIX V4.1. However, sometimes you might have some troubles. This section summarizes several common problems and indicates possible solutions.

3.7.7.1 Problem Determination

Basically, you should determine whether the problem is in the SNA API or not. If not, you should check on commands or subroutines which have changed in AIX V4.1. Check to see if your SNA API call returns an error. Looking into the file /usr/include/luxsna.h for the SNA error codes can also give you some hints as to what is wrong in your Transaction Programs. Conversation flow is also important. These are typical examples of possible errors:

Init Failure: An error in the `cminit` or `snaopen` subroutines shows a bad Initialize Conversation parameter. Check your Side Information profile, Local LU profile and the name of the Transaction Profile Name (TPN) profile. Particularly after using the `migratesna` command or adding new profiles, it is likely that you may be using the wrong profile name.

Alloc Failure: An error in the `cmalloc` or `snalloc` routines shows an Allocate Conversation failure. First, check the session state with the command:

```
sna -d s
```

If no session exists, check the session profile parameters—especially the Partner LU profiles to see that the Partner LU name is correct and that the Side Information profile and the Session Security profile are correct.

If the session state is good, check your Max Session value in the Mode profile. Also, check the conversation type or sync-level in the Transaction Program Name profile or your Transaction Programs. If you use multiple instances of your transactions, change the multiple instances parameter in the TPN profile to **yes** (the default is **no**). In the case of migration, those errors may be caused by a profile name mismatch.

First Send Failure: Communications Server Version 4 will not immediately send an `allocate` command to the remote side. Data will be sent at the time the buffer is filled or when an SNA command with a buffer flush, such as `cm_confirm` or `cm_flush`, is issued. If your send command returns with an error, it might be a problem with the Remote Transaction Program. Check the Remote Transaction Program Name in the Side Information profile, and the Transaction Program Name profile on the remote side. Also check your file permissions. Sometimes, the user named within the TPN profile is not allowed to execute this Transaction Program, or the path given is not the actual path to the executable code.

3.7.7.2 API Trace

The API trace is also very useful. You can request the following types of API traces:

- Operating System Subroutines and Library Subroutines API trace (such as `snaopen`, `snactl`)
- CPI-C API trace (such as `cmalloc`, `cmsend`)
- Generic SNA API trace (such as `open`, `write`)
- LU 0 SNA API trace

Perform the following steps to start, stop, and format the API trace:

1. Use the following command to start the trace:

   ```
   sna -trace -c on
   ```

2. When the Transaction Program has finished executing, use the following command to stop the trace:

   ```
   sna -trace -c off
   ```

3. To close the active trace file and enable it for formatting, issue the following command:

   ```
   sna -setlogs -t
   ```

4. Use the following command to format the trace information:

   ```
   trcrpt -d API_type_number /var/sna/snaservice.n
   ```

   where `API_type_number` is:

   271 Operating System Subroutines and Library Subroutines API

   390 CPI-C API

   281 Generic SNA API

   In this example, `n` is the suffix of the log file containing the trace data.

Starting and Stopping an LU 0 API Trace: You can also generate an API trace for an LU 0 primary or secondary LU by selecting **yes** for the Enable API trace field in the LU 0 Primary LU profile or in the LU 0 Secondary LU profile. The trace begins when a session for the LU starts and ends when the session ends.

(You can end the LU 0 session by including the `lu0close` or `lu0closep` subroutines in your program.)

The LU 0 API traces are named /var/lu0/PrimaryLUProfileName.pri_api and /var/lu0/SecondaryLUProfileName.pri_api.

3.8 Migrating HCON

This section explains the steps necessary for migrating the HCON application.

3.8.1 HCON Administrator Definitions

Independent of the overall installation method used, the new HCON Version 2.1 must be installed separately. The installation process will overwrite the file /etc/hcon/users which holds all information about users that are defined to HCON. This is known as the HCON Administrator definitions. If you only have one or two users, this might not cause any problems. But if you have hundreds of users, this will cost you hours to rebuild all your HCON definitions. Therefore, hopefully, you performed the recommended actions and saved the old definitions. You now have two options:

1. Start from the beginning by adding each HCON user.
2. Restore your backup to /etc/hcon/users.

If you choose the first option, when adding every user, all old HCON definitions that may still exist in that user′s home directory will be overwritten.

In the second option, nothing happens, except HCON now knows all the previous users. You must also restore the HCON definitions in the user′s HOME directories in order to use HCON.

3.8.2 HCON User Definitions

Every HCON user has their own definitions, including the number of sessions that can be used and the keyboard mappings and colors to use for each session. These files are usually stored in the user′s home directory:

- $HOME/usrprofs
- $HOME/usrdflts
- $HOME/SYS*
- $HOME/e789_ktbl

  This file, if present, can actually use any name. The name is listed in the users′s session definition.

- $HOME/e789_ctbl

  Likewise, this optional file can use any name.

Depending on the type of the installation, these definitions may be still in place. They will be destroyed by adding or removing the specific HCON user.

Depending on all the actions taken before, you can use the old definitions either by leaving them untouched or by restoring them from your backup media.

HCON sessions can be used on different link types:

- TCP/IP link

- Coaxial cable
- SNA link

For TCP/IP or coaxial connections, nothing has changed after the migration, and you can use HCON immediately without any change. In the case of SNA connections, you may have to make some modifications:

- If you are migrating from SNA Server/6000 Version 2.1, you can continue to use HCON immediately.
- If you are migrating from an SNA Services/6000 Version 1.2.1 environment, you should check your HCON session definitions.

HCON on SNA Services/6000 Version 1.2.1 systems points to *Connection* profiles. These profiles are retired in Communications Server Version 4 and are now called LU 1, 2, and 3 *Session* profiles. In most cases, the new Session profile will have the same name as the old Connection profile; so your HCON definitions should still be correct. However, if you have problems, check the HCON profile, and look for the right name for the LU 1, 2, or 3 session profile.

Note: The field name still refers to the **SNA logical connection**.

```
                        Reconfigure SNA Display Session

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

[TOP]                                                    [Entry Fields]
  HCON user name                                        root
  Session name                                          a
  Session USE                                          [snasess02]
* SNA logical connection prefix or profile             [con202]
  LANGUAGE                                              English (U.S.A.)
* KEYBOARD table                                       [/usr/lib/hcon/e789_
* COLOR table                                          [/usr/lib/hcon/e789_
* File used by SAVES key                               [/e789_saves]
* File used by REPLS key                               [/e789_repls]
* Local printer used by PRINT key                     [lp0]
  Host TYPE                                             CMS
  Host LOGIN ID                                        []
  Autolog NODE ID                                      []
  Autolog TRACE                                         no
[MORE...6]

F1=Help            F2=Refresh          F3=Cancel           F4=List
```

3.9 Migrate Other AIX-SNA Products

The SNA family consists of several products. General migration rules for some of the other products are described below.

3.9.1 SNA Application Access for AIX

This product allows access from any SNA device in an SNA network to an application running on AIX. The best example of this is the access to CICS/6000 from any 3270 device in the network. Since SNA Application Access uses its own networking routines, it generally does not use Communications Server Version 4 functions.

Customers using this product on AIX Version 3.2 systems cannot move to AIX Version 4.1 because, at the time this document was written, SNA Application Access was not supported on AIX Version 4.1.

3.9.2 SNA Client Access Version 1.1

SNA Client Access for AIX allows any TCP/IP client to connect to an SNA network by using the tn3270 protocol. Version 1.1 runs on both AIX Version 3.2 and AIX Version 4.1 without any change.

Since product and configuration files are stored in the /usr/lpp/SNA_CA directory, you should be aware that these files will be deleted by the Complete Overwrite and Preservation install methods. Therefore, you should at least save all of your configuration files *my_name.cfg* found within different subdirectories of /usr/lpp/SNA_CA.

This product also depends upon the definitions of your SNA configuration. These definitions are saved in the SNA Server/6000 profiles and will be saved and restored when following the Communications Server Version 4 migration steps.

After migrating the AIX system, you have to re-install the SNA Client Access product, and restore the old *.cfg files. When you complete migrating your SNA configuration, SNA Client Access should operate as before.

3.9.3 SNA Client Access Version 1.2

This new release of SNA Client Access extends the function by using the tn3270e protocol. It allows you to use 3287 printer datastream over a TCP/IP network. While Version 1.1 could only use pooled LUs, the Version 1.2 allows the selection of specific SNA LUs in order to connect a client application to a predefined SNA LU. This function is important for printer emulations where you have to know the exact LU.

SNA Client Access Version 1.2 will be available in February and March 1996 for both AIX Version 3.2 and AIX Version 4.1.

Note: The SNA Client Access Version 1.2 for AIX V3.2 and the SNA Client Access Version 1.2 for AIX V4.1 are different products with different installation images. As part of planning for your AIX upgrade, you must order a refresh of the product to obtain the AIX Version 4.1-compatible code.

The product files for Version 1.2 are still stored in /usr/lpp/SNA_CA, but the configuration files are now stored in the /etc/aixsnaca directory. After migration of the AIX system, and installation of the new SNA Client Access Version 1.2 for AIX Version 4.1 code, you can use the configuration files from either SNA Client Access V1.1 or V1.2. However, SNA Client Access V1.1 configuration files must be moved to the /etc/aixsnaca directory.

Chapter 4. Sample SNA Migration

This section shows a typical migration example. In reality, your system could be simpler or more complex, but the migration path will be basically the same. This sample demonstrates an outline of the migration.

4.1 Test Environment

The sample system environment is shown in Figure 8.

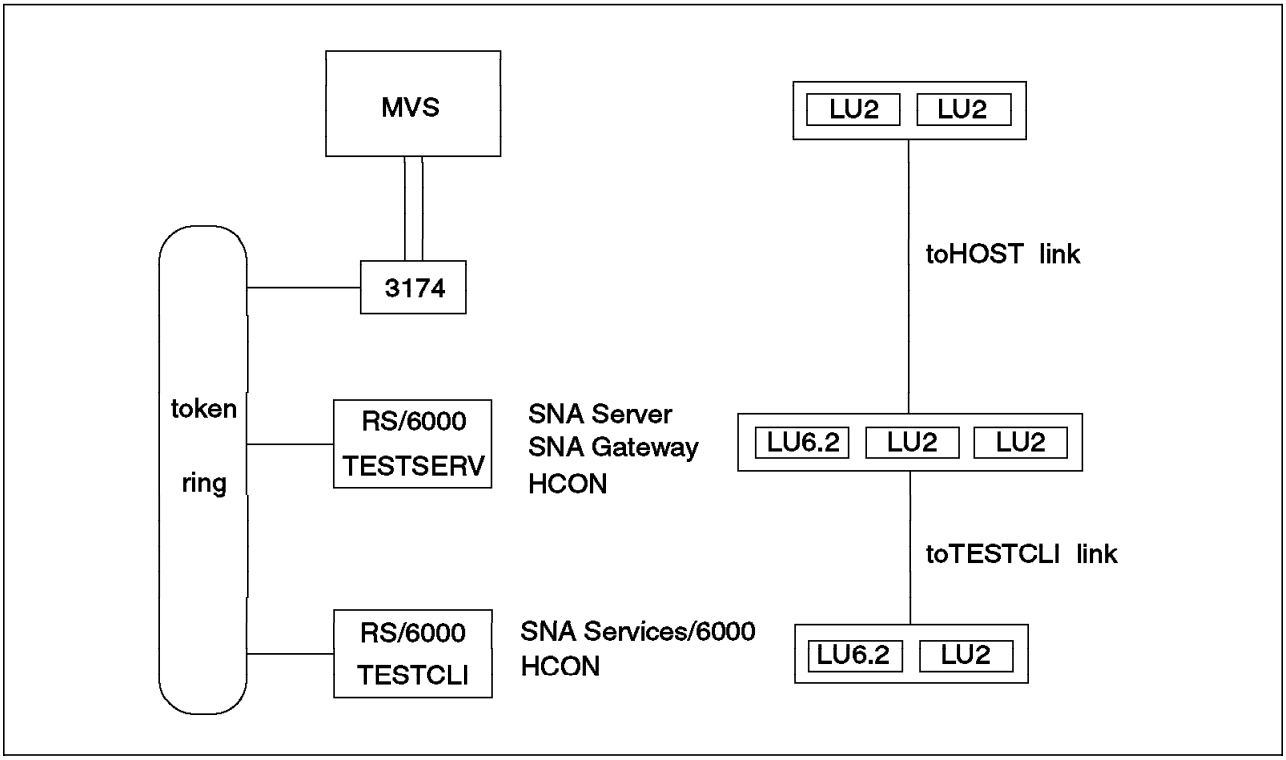

Figure 8. Migration Sample Environment

The levels of operating-system and SNA software before the migration are shown in Table 12.

Table 12. Sample Environment

System Name	TESTSERV	TESTCLI
Operating System	AIX Version 3.2.5	AIX Version 3.2.5
LPPs	SNA Server/6000 V2.1	SNA Services/6000 V1.2
	SNA Gateway/6000 V2.1	
	HCON Version 1.3	HCON Version 1.3

Our sample systems were connected via a token-ring to an IBM 3174 to provide access to an MVS host. One RS/6000, named TESTSERV, had four sessions (LU type 2) to the MVS system. Two of these sessions were used for HCON directly on TESTSERV. The TESTSERV system also ran the SNA Gateway/6000 software which it used to provide the other two LU2 sessions to the workstation, named TESTCLI. Additionally, the two RS/6000s shared four direct LU type 6.2 sessions

for APPC communication. Each RS/6000 had four Transaction Programs for file transfer and for sending messages. The test Transaction Programs were the following:

CPICSEND This program uses CPI communications to send a file, and includes data conversion from ASCII to EBCDIC.

CPICRCV This program uses CPI communications to receive a file, and includes data conversion from EBCDIC to ASCII.

SUBRSEND This sends messages to SUBRRECV using SNA Library Subroutines.

SUBRREC This receives messages from SUBRSEND using SNA Library Subroutines.

Both systems were migrated to AIX V4.1 and Communications Server Version 4. TESTSERV was done using the Preservation installation method, while the Complete Overwrite installation was used on TESTCLI.

Note

We did not use the Bundle Install for Communications Server Version 4 which automatically installs SNA, Gateway, AnyNet, and HCON functions.

4.1.1 TESTCLI Migration (Complete Overwrite)

This system had two LU type 2 session definitions used for HCON terminal emulation, and one LU type 6.2 definition for Transaction Programs. This single LU 6.2 session was used for four separate sessions. Since the system had SNA Services/6000 Version 1.2.1 installed, only one link could be defined and active for all SNA sessions. Therefore, we had one Attachment profile and three Connection profiles. The profiles and definitions used can be found in A.2, “Old Profiles on TESTCLI” on page 271. In order to migrate to AIX 4.1 and Communications Server Version 4, we used the Complete Overwrite installation method.

On the following pages, step-by-step instructions are provided to show how this migration was done.

4.1.1.1 Backup

To prepare for the installation, we first backed up the SNA profiles, Transaction Programs and user data. The backup steps are detailed below:

1. Exporting the SNA profiles:

   ```
   exportsna -f /home/test/test.pro.old
   ```

2. Saving the HCON user definitions:

   ```
   cp /etc/hcon/users /home/test/hcon.users
   ```

3. Saving the HCON definitions from the user’s HOME directory:

   ```
   cd /home
   tar -cvf /home/test/hconusr.tar find . -name usr* -print | xargs
   tar -cvf /home/test/hconsys.tar find . -name SYS* -print | xargs
   ```

4. Saving the user applications:

```
cd /usr/lpp/sna/samples/bin
tar -cvf /home/test/appl.tar CPIC* SUBR*
```

5. Backing up all saved information to a backup media (tape).

```
cd /home/test
tar -cvf /dev/rmt0 .
```

Note: This is just a sample of a backup procedure—you may use other methods and other media, and you may have more or less data and files to save.

4.1.1.2 Installing AIX V4.1

We were then ready to install AIX V4.1 using a Complete Overwrite installation. The following file sets were the installed:

AIX V4.1.3 Client Bundle	(includes bos.rte.iconv)
bos.dlc.com	Common Data Link Control files
bos.dlc.token	Token-Ring Data Link Control
bos.iconv.com	Common Language to Language
bos.iconv.ja_JP	EBCDIC and ASCII Language
bos.sysmgt.trace	Software Trace Service Aids

After installation, disk space used was 224 MB (this included 64 MB paging space and 152 MB for the /usr filesystem).

4.1.1.3 Installing Communications Server Version 4

Only the SNA Server Base filesets (LU1, LU2, LU3, and L6.2) were selected, which automatically included the messages:

sna.rte	SNA Server Base (LU1 LU2 LU3 L6.2)
sna.msg.en_US	SNA Server Messages

Be careful when selecting the filesets with the SMIT install menu. If you choose **sna ALL**, then all SNA components will be installed automatically, which will take a lot of disk space. The online documentation filesets, *sna.books.**, were not installed. After the installation, the disk space used was 264 MB (/usr was extended to 168 MB and /tmp to 12 MB).

4.1.1.4 Installing HCON V2.1

You do not have to install all HCON file sets. In this sample system, only these filesets were installed:

hcon.rte	3270 Host Connection Program
hcon.util	HCON Utility Program
hcon.X11	HCON for AIXWindows
hcon.msg.En_US	HCON Messages
hcon.terminfo.ibm.data	HCON Terminal Definitions - IBM

The final disk space usage was 268 MB, with /usr being extended to 176 MB.

After completing the installation, we restored all the user data, Transaction Programs, SNA profiles, and HCON profiles previously saved to tape. This is shown in the next steps.

4.1.1.5 Customizing the System

Since a Complete Overwrite installation was performed, TCP/IP had to be customized from scratch, and all the users had to be added to the system.

1. We made the TCP/IP interface by using the command:

   ```
   /usr/sbin/mktcpip -h'TESTCLI' -a'9.3.1.39' -m'255.255.255.0' -i'tr0' -r'16'
   ```

2. We then added the user definitions:

   ```
   mkuser 'old_user_name'
   passwd 'old_user_name'
   ```

Note: Do not forget to set the password for every user in order to enable logins.

4.1.1.6 Restoring the Saved Information

We then restored our SNA definitions from the backup tape:

1. Making a directory:

   ```
   mkdir /home/test
   ```

2. Restoring the tape files:

   ```
   cd /home/test
   tar -xvf /dev/rmt0
   ```

3. Restoring user applications CPICSEND/RECEIVE and SUBRSEND/RECEIVE:

   ```
   cd /usr/lpp/sna/samples/bin
   tar -xvf /home/test/appl.tar
   ```

4. Restoring the HCON definitions:

   ```
   cp /home/test/hcon.users /etc/hcon/users
   cd /home
   tar -xvf /home/test/hconusr.tar
   tar -xvf /home/test/hconsys.tar
   ```

4.1.1.7 Migrating SNA profiles

The next step involved the migration of the SNA profiles into the format required for the new SNA version by entering the following commands on the AIX command line:

```
migratesna -s /home/test/test.pro.old -t /home/test/test.pro.new \
    -e error_message_file -l NET1.CP2
```

We could instead have used the SMIT menu interface:

```
smit migratesna
```

The full Local Control Point name, for example NETID.CP_NAME, is required for this command. This step generates the profiles required for the new SNA version. In our case, there were no severe errors during the migration step.

4.1.1.8 Importing and Verifying the SNA profiles

Next we imported the new SNA profiles into Communications Server Version 4.

```
importsna -f /home/test/test.pro.new
```

In this migration example, we didn't get any errors, so we proceeded to the step of verifying and updating the SNA database.

```
verifysna -U
```

You may see some errors or warnings during this process. In our case, we again had no errors. You should read all messages and warnings carefully, and you should correct any errors until you achieve an error-free verification step.

4.1.1.9 Migrating Transaction Programs

In this example case, all the Transaction Programs and HCON sessions worked well after the migration. The only necessary changes were in the shell scripts used for starting the Transaction Programs, which were modified to use the `sna` command instead of using `startsrc` commands. All four Transaction Programs operated correctly without the need for any modification. If your Transaction Programs are well writen and well behaved, it is likely that your migration will be as easy as this sample.

This is a list of transactions used in the sample Transaction Programs:

cminit	Initialize Conversation
cmssl	Set Sync_Level
cmsrt	Set Receive_Type
cmallc	Allocate Conversation
cmsdt	Set Deallocate_Type
cmsst	Set Send_Type
cmsend	Send Data
cmdael	Deallocate Conversation
cmaccp	Accept Conversation
cmrcv	Receive Data
cmcfm	Request Confirm
cmcfmd	Response Confirm
snaopen	Initialize Session Parameters
snalloc	Allocate Conversation
snactl	Control Conversation
snaread	Receive Data
snawrit	Send Data
snadeal	Deallocate Conversation
snaclse	Free the Indicator

These sample Transaction Programs did not use any of the structures where fields have changed between versions. Use of these structures would increase the migration work required.

4.1.1.10 Post-Installation Procedures

To complete the installation, we continued with post-installation procedures:

1. iFOR/LS keys

   Since we were installing SNA from a full distribution media, the required iFOR/LS key was installed automatically.

2. Session count

We set the session count field to the required level. In a customer site, this would correspond to the licensed number of sessions. The session count can be set by using the command:

```
chsnaobj -t'sna' -a '400' sna
```

or by using

```
SMIT _snasna
```

```
                        Change/Show SNA Node Profile

Type or select values in entry fields.
Press Enter AFTER making all desired changes.
                                                        [Entry Fields]
  Profile name                                          sna
  Maximum number of licensed sessions (1-5000)         [400]
  Maximum number of conversations (1-5000)             [200]
  Restart action                                        once
  Dynamic inbound partner LU definitions allowed?       yes
  NMVT action when no NMVT process                      reject
  Trusted group names                                  [system]
  APPC security sense codes                             specific
  Start SNMP subagent when SNA is started?              no
  Time out limited resource sessions?                   no
    If yes, time-out value (1-3600 seconds)            [15]
#
```

3. Changing the Log Size

   We changed the log to hold only 500 KB of data and to use only three logs files. We used SMIT _snasetlogs to change these values:

```
                      Configure SNA Log and Trace Files

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

                                                        [Entry Fields]
  Set file names and sizes
    Use SNA or System file as service log?              SNA
      Service log file name                            [/var/sna/snaservice
      Service log file size                            [500000]
    SNA event/flow trace log file name                 [/var/sna/snatrace]
    SNA event/flow trace log file size                 [0]
    SNA failure log file name                          [/var/sna/snafailure
    SNA failure log file size                          [0]

    File wrap limit                                    [3]
```

4. We also uncommented the SNA start command, sna -start , in the /etc/rc.sna file to automatically start SNA every time the system is rebooted.

4.1.1.11 Connection and Application Test

In order to verify the successful migration, we then started SNA and the SNA applications.

1. While logged in as the root user, we started SNA with the command sna -start.
2. We then logged in as an HCON user and started a terminal emulation session by typing e789 a. This started the Link Station to the host system, and the VTAM login screen appeared.
3. We then tested all of our applications for correct operation.

In our case, all was working fine; so we could start the migration of our server system.

4.1.2 TESTSERV Migration (Preservation)

Before migration, the TESTSERV system was an AIX V3.2.5 system with SNA Server/6000 Version 2.1 installed. This system had four LU type 2 session definitions and one LU type 6.2 definition. Two LU type 2 sessions were used locally; the other two sessions were supplied for the downstream workstation (TESTCLI) through SNA Gateway/6000. The single LU 6.2 definition was again used for four concurrent sessions. TESTSERV had one SNA DLC profile and two link station profiles with different XIDs. The two Link Stations were *toHOST* and *toTESTCL*. The same Transaction Programs as on the TESTCLI system were used for file and data transfer.

We chose to use the Preservation installation method to upgrade this system.

4.1.2.1 Backup

To get ready for installation, we again backed up the SNA profiles, TPs and user data. This was done in the same manner as on the system TESTCLI. See 4.1.1.1, “Backup” on page 104 for the actual commands used.

1. Exporting the SNA profiles was done with the command:

   ```
   exportsna -f /home/test/server.pro.save
   ```

The backup step also included

2. Backing up the HCON users file:

   /etc/hcon/users

3. Backing up the user applications and data (including HCON profiles in the users′ home directories).

Actually, in a Preservation installation, user data in the /home directory will be preserved. However, backing up your data is always a great idea, just in case of accidents.

4.1.2.2 Installing AIX V4.1

After backing up profiles and user data, we installed AIX V4.1 with the Preservation installation method. We installed the following filesets:

AIX V4.1.3 Server Bundle	(includes bos.rte.iconv and bos.sysmgt.trace)
bos.dlc.com	Common Data Link Control files
bos.dlc.token	Token-Ring Data Link Control
bos.iconv.com	Common Language to Language
bos.iconv.ja_JP	EBCDIC and ASCII Language

Note: If you use the Migration installation method, DLC and iconv filesets will be installed automatically.

4.1.2.3 Installing the SNA Related Software

We did not use the Bundle installation which would automatically install all three products.

These filesets were installed for SNA and HCON: (messages are automatically installed for each base product.)

sna.rte	SNA Server Base (LU1 LU2 LU3 LU6.2)
sna.msg.en_US.rte	SNA Base Messages - en_US
gw.rte	SNA Gateway
hcon.rte	3270 Host Connection Program
hcon.util	HCON Utility Program
hcon.X11	HCON for AIX Windows
hcon.msg.En_US	HCON Messages
hcon.terminfo	HCON Terminal Definitions ALL

4.1.2.4 Restoring the Saved Information

After completing the installation, we had only to restore the SNA profiles and HCON users definition file. HCON profiles, SNA Transaction Programs and user data were preserved as they were stored in the /home filesystem. A Migration installation will save all definitions, but because of the new installation of HCON, the /etc/hcon/users file will be overwritten and must be restored anyway. So, the restoration has only one step:

1. `cp /home/test/hcon.users /etc/hcon/users`

4.1.2.5 Importing and Verifying the SNA Profiles

Next we imported the SNA profiles from /home/test/server.pro.save.

1. We import the saved profiles by using the command:

   ```
   importsna -f /home/test/server.pro.save
   ```

   We could also have used the SMIT interface:

   ```
   smit importsna
   ```

2. The profiles were verified using the commands:

   ```
   verifysna -U
   ```

   Again, we could equally have used:

   ```
   smit verifysna
   ```

If you use a Migration installation, you need only perform the `verifysna` command after installation; the back up and import process can be skipped. SNA migration from SNA Server/6000 V2.1 on AIX V3.2 is simple; there are no changes to be made in the profiles.

4.1.2.6 Confirmation of Transaction Programs

Next we were ready to test our Transaction Programs. This sample case worked with no problems. No profiles or Transaction Programs had to be modified. In the case of migration from SNA Server V2.1, you need only watch that you install all required filesets, and be aware of possible AIX command changes.

4.1.2.7 Post-Installation Procedures

Again, we completed the installation by following these steps:

1. Checking the iFOR/LS keys
2. Setting the session count according the license agreement
3. Setting the log sizes and counts
4. Having SNA start automatically on reboots by uncommenting the lines in /etc/rc.sna

4.1.2.8 System Test

In order to verify the updgrade, we then started SNA and all Link Stations, and tested the HCON and Gateway functions.

1. While logged in as the root user, we started SNA with the command:

   ```
   sna -start
   ```

2. We then started the Link Stations:

   ```
   sna -start l -p toHOST
   sna -start l -p toTESTCL
   ```

We received the messages:

```
0105-2723 The "toHOST" Link Station has been started
```

In order to see the status of your links, you can use the command `sna -d l`.

```
    Link           Adjacent      Node    Device            # of local  In
   station         CP name       type     name      State    sessions   use
-------------- ----------------- ----- --------- ---------- ---------- ----
@tok0.4                               tok0      Starting            0 No
toTESTCL                              tok0      Starting            0 No
toHOST         USIBMSC.SCG20     LEN   tok0      Active              2 Yes
```

The Link Station toTESTCL is a listening Link Station and will become active when a user on the TESTCLI system starts an HCON session or when the root user starts the Link Station to our TESTSERV system.

The activated LUs can be displayed by the command:

```
sna -d s123
```

```
 Local  LU  Session    Session         Local   Host       Link
  ID   type  type     profile name    LU addr SSCP ID    station       Status
------ ---- -------- ---------------- ------- ------- -------------- --------
2      2    SSCP-LU  toHOST03         3       5411    toHOST         Active
1      2    SSCP-LU  toHOST02         2       5411    toHOST         Active
#
```

The activated LUs for the Gateway can be seen by using the command:

```
sna -d gw -l toHOST -o l
```

```
************************************
        Gateway Information
************************************
Number of gateway Link Stations         1

1>Configured Link Station               toHOST
  Active Link Station                   toHOST
  Node Type                             Host
  Control Point Name                    USIBMSC.SCG20
  State                                 Active
  Maximum BTU size on the link          2052
  Total gateway LUs on this link        2

1.1>LU Address                          4 (0x04)
    State                               Active - LU Disabled

1.2>LU Address                          5 (0x05)
    State                               Active - LU Disabled
```

After successfully starting the Link Stations, you should be able to use all HCON sessions and all your Transaction Programs as before. This was the last step in our migration sample. The SNA definitions used, both before and after the upgrade, can be found in Appendix A, “SNA Definitions and Profiles” on page 271.

Chapter 5. Migration of Special Link Types

The SNA migration considered so far covers those link types where support is included in the base operating system and base SNA code. However, there are several other types of links that are commonly used in AIX communications that involve separate device drivers or device support software that is not included in the base code. The migration of systems using some alternative link types is described in the following sections.

Note: Although this section is concerned mostly with the migration of these link types when used in an SNA environment, much of the chapter applies equally to the use of the X.25 networks and channel interfaces for TCP/IP transmission.

5.1 ESCON and BLKMUX Channel Connectivity for AIX

One common requirement in SNA networks is to communicate with a mainframe system over a channel architecture to provide very high bandwidth data transfer rates. In the RISC System/6000, this can involve the following two adapters:

- Block Multiplexer (BLKMUX) Channel Adapter
- Enterprise System Connection (ESCON) Control Unit Adapter

In AIX Version 3.2.5, support for these adapters came from either the Channel Connectivity features of AIX or from the SNA Channel Connectivity features of SNA Server/6000 Version 2.1. Neither of these sources can be used after migrating to AIX Version 4.1. Under AIX Version 4.1, support for these adapters comes in the form of separate program products. The following sections give you the information necessary to migrate any of these products.

5.1.1 Channel Migration Path

The migration of any channel-to-host connection can be shown in simple steps. Detailed information on the important steps will start with 5.1.2, “Channel Adapter Microcode” on page 114.

1. Document your channel device definitions
2. Document your channel TCP/IP definitions
3. Document your channel CLIO/S definitions
4. Export your channel SNA definitions, and save them on backup media
5. Get the newest microcode
6. Get the new channel support program
7. Migrate the AIX system to AIX Version 4.1
8. Install the new microcode
9. Install the new channel support program
10. Install the new version of SNA and the SNA channel support feature
11. Install APAR PN74395 for CLIO/S if applicable
12. Re-enter or check the channel device definitions
13. Re-enter or check the channel TCP/IP definitions
14. Re-enter or check the channel CLIO/S definitions

15. Migrate the SNA definitions

16. Test the system

5.1.2 Channel Adapter Microcode

In addition to the device support; to use the channel adapters, you also require microcode that will be downloaded to the adapter. The channel adapter microcode, and the method of obtaining that microcode, have changed several times. You must have the appropriate microcode level in order to make the channel adapter available. The list below gives you a summary of the different device support programs (device drivers DD) and their corresponding microcode sources.

AIX V3.2.5, AIX BLKMUX Device Driver
The microcode came on diskette with the adapter.

AIX V3.2.5, AIX ESCON Device Driver
The microcode came on diskette with the adapter.

AIX V3.2.5, SNA BLKMUX Device Driver
The microcode was included in the SNA device driver fileset sna.blkmux.cuu.

AIX V3.2.5, SNA ESCON Device Driver
The microcode was included in the SNA device driver fileset sna.escon.cuu.

AIX V4.1, Channel Support BLKMUX Device Driver
A microcode Engineering Change (EC) is needed to update your adapter for AIX V4.1. EC number D26567A, part number 39H9227 can be ordered through your hardware support channel. IBM staff can obtain BLKMUX PACKAGE from the AIXTOOLS disk, which contains the Licensed Program Product (LPP) blkmux.mc at level 3.2.0.10.

AIX V4.1, Channel Support ESCON Device Driver
A microcode EC is again needed. EC number C74216E, part number 40H2902 can be ordered through your hardware support channel. IBM staff can obtain ESCON PACKAGE from the AIXTOOLS disk, which contains the LPP escon.cuu at level 3.2.0.1.

The microcode is not migrated; the new microcode level comes as an installp image and is simply reinstalled. If you order a new channel adapter with an AIX Version 4.1 system, the new adapter microcode will be shipped on diskette with the adapter.

Note

If you are migrating an existing AIX Version 3.2 system with a channel adapter, you must be aware that you need to order the microcode EC to obtain the correct microcode levels. IBM staff can also download the microcode from the AIXTOOLS disk.

The microcode is backwards compatible. The newest microcode will support all of the device support programs.

5.1.3 Channel Device Driver Migration

The steps required for the migration of your channel device driver software will depend upon the software being migrated from. In some cases there are no utilities to export and import or migrate your configuration, so you will have to reconfigure your channel from scratch. The different cases are shown in Table 13 below.

Table 13. Channel Device Driver Migration

Old Driver		New Driver		Steps Req'd
5750-037	AIX BLKMUX Device Driver	5765-604	AIX BLKMUX Channel Connectivity	**a**
5750-037	AIX ESCON Device Driver	5765-603	AIX ESCON Channel Connectivity	**b**
5765-247	SNA BLKMUX Feature	5765-604	AIX BLKMUX Channel Connectivity	**c**
5765-247	SNA ESCON Feature	5765-603	AIX ESCON Channel Connectivity	**d**

a Migration is completely manual. The definitions must be re-entered because the format and SMIT screens are completely different.

b, c, d The steps required in these cases depend on method used to upgrade your system:

- Preservation or Overwrite

  Manually re-enter all definitions

- Migration Install

  In these cases, the migration process automatically preserves most of your configuration, but some changes must be made:

  – The *Connection* definition no longer exists. It is automatically created, changed or deleted when the associated *SubChannel* definition is created, changed or deleted. This makes the configuration easier for the user.

  – TCP/IP naming restrictions. The TCP/IP SubChannel definition, supported only in Common Link Access to Workstations (CLAW) mode, must be named the same as the corresponding TCP/IP network interface. For block multiplexer channels, this is ca0, ca1 and so on. For ESCON channels, use es0, es1 and so on.

5.1.4 Channel - SNA Definition Migration

For migration of SNA profiles, use the normal method. Note that we need only consider the case of migration from SNA Server/6000 Version 2.1, as there was no support for channel communications in SNA Services/6000 Version 1.2.1.

1. `exportsna -U -A -f filename`
2. Save the exported file to /home or other safe place. Remember that depending on the installation method, /home may not be safe.
3. Upgrade to AIX Version 4.1
4. Install Communications Server Version 4
5. `importsna -f filename`

   You may see the message:

```
0105-0393 Unrecognized field 'rrm_enabled' in profile type 'sna'
```

This message appears because the rrm_enabled field has been moved from the SNA Node Profile to one of the lu62 profiles. You can ignore this message.

6. verifysna -U

5.1.5 Channel Product Installation

The following example shows the installation of the AIX ESCON Channel Connectivity program, feature code 5765-603. The steps for installing the BLKMUX software are identical.

The AIX ESCON Channel Connectivity program contains the device drivers for the adapter and includes the TCP/IP support. The SNA channel support feature is then installed on top of the device support and installs the data link control (DLC) necessary for SNA access to the channel device driver.

- To install the new product, with the adapter installed in the system, use the command:

  cfgmgr -i <install device>

  Where <install device> is the name of the device that contains your device driver software. You can also use the SMIT menus by typing:

  smit cfgmgr

  You will see the following screen:

```
                    Install/Configure Devices Added After IPL

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

                                                        [Entry Fields]
  NOTE: A selection of "none" configures devices
  added after IPL without the installation of
  software
  INPUT device / directory for software              [/dev/rmt0]             +









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

The successful installation of the AIX ESCON Channel Connectivity program gives the following output:

```
Installation Summary
--------------------
Name                        Level           Part        Event       Result
-------------------------------------------------------------------------------

devices.mca.8fc3.diag       4.1.4.0         USR         APPLY       SUCCESS
devices.mca.8fc3.rte        1.1.0.0         USR         APPLY       SUCCESS
devices.mca.8fc3.rte        1.1.0.0         ROOT        APPLY       SUCCESS

************************************************************
* The ESCON lpp has been  installed on your machine.      *
*                                                          *
* After adding any ESCON  device driver definitions,      *
* you can configure your  ESCON adapter by running        *
* "configuring devices after ipl" from the devices menu   *
* in SMIT.  To add ESCON  device driver definitions       *
* run "smitty esca".                                       *
************************************************************
```

- Install Communications Server Version 4 and the SNA Channel Support feature using the *Install Software Products at Latest Level* path on SMIT.

```
Installation Summary
--------------------
Name                        Level           Part        Event       Result
-------------------------------------------------------------------------------

sna.rte                     3.1.1.0         USR         APPLY       SUCCESS
sna.dlcchannel              3.1.0.0         ROOT        APPLY       SUCCESS
```

For samples of ESCON or BLKMUX definitions, you should refer to *RISC/6000 to Mainframe Using S/370 Channel Connections*, SG24-4589.

5.2 Migrating X.25

In AIX Version 3.2, a base level of X.25 support was provided with the operating system. In addition, the Licensed Program Product (LPP) AIX X.25 Version 1.1 (5696-868) was later introduced to extend the X.25 support with additional features. In AIX Version 4.1, there is no X.25 support in the base operating system. The LPP AIXlink/X.25 Version 1.1 (5696-926) provides AIX Version 4.1 with identical functionality to the AIX X.25 Version 1.1 LPP on AIX Version 3.2. AIXlink/X.25 Version 1.1 is fully backwards compatible with the X.25 support provided with AIX Version 3.2, and is therefore especially appropriate for customers currently using X.25 on AIX Version 3.2-based systems who wish to migrate to AIX Version 4.

For a full description of the X.25 LPP features and enhancements, and of X.25 migration, we recommend the following documentation:

- *AIXlink/X.25 1.1 for AIX, Guide and Reference*, SC23-2520
- *AIX/6000 X.25 LPP Cookbook*, GG24-4475

We will give a brief outline of some of the changes and migration steps below.

5.2.1 X.25 Differences Between AIX Versions

The following table shows the main characteristics and differences between X.25 support on AIX Version 3.2 and AIX Version 4.

Table 14. X.25 Differences Between AIX Version 3.2 and AIX Version 4.1

	AIX V3.2 BOS X.25 Support (included in 5756-030)	**AIX X.25 V1.1 (5696-868)**	**AIXlink/X.25 V1.1 (5696-926)**
AIX Operating System	AIX Version 3.2	AIX Version 3.2.5 or above	AIX Version 4.1.1 or above
Installation	Single installable image	Single installable image	Selectable filesets
Adapter microcode	Shipped on diskette with adapter	Installed with LPP	Installed with LPP
X.25 Co-processor Adapter name	x25s0 (AIX levels <3.2.3e) ampx0 (AIX levels 3.2.3e and above)	ampx0	ampx0
X.25 Portmaster Adapter name	Not supported	apm0	apm0
Driver name	x25s0	twd0 (streams device)	twd0 (streams device)
Port name	x25s0	sx25a0	sx25a0

Some of the new features supported by the X.25 LPP include:

- Support for the International Telegraph and Telephone Consultative Committee (CCITT) 1988 X.25 standard
- Packet layer programming interfaces
  - Network Provider Interface (NPI)
  - Data Link Provider Interface (DLPI)
- Packet Assembler/Disassembler (PAD) support
- V.25bis support
- Automatic DTE configuration

There are some other major changes in handling and controlling X.25 on the RISC System/6000.

Note: These changes apply when moving from the AIX BOS X.25 support on AIX Version 3.2 to AIXlink/X.25 support on AIX Version 4.1. If you were using the AIX X.25 LPP on AIX Version 3.2, you will notice very few differences in X.25 support when upgrading to AIX Version 4.1.

xmanage This command could be used in the old version to activate or deactivate the X.25 interface and display the status. The command `xmanage` is no longer available. Activation or restart and recovery will be done automatically, depending on device parameters or applications.

xmonitor This command was used to start a link level trace. A new command `x25mon` is provided with the new product in order to start any kind of link trace for X.25.

xroute This command is still available in order to provide a routing of incoming calls to the appropriate application. However, this routing is only valid for applications using the Common Input/Output (COMIO) interface, such as SNA or xtalk. TCP/IP no longer uses this table.

Figure 9 should help you to understand the hierarchical structure of the X.25 product for AIX Version 4.1. Some features of the diagram are explained below.

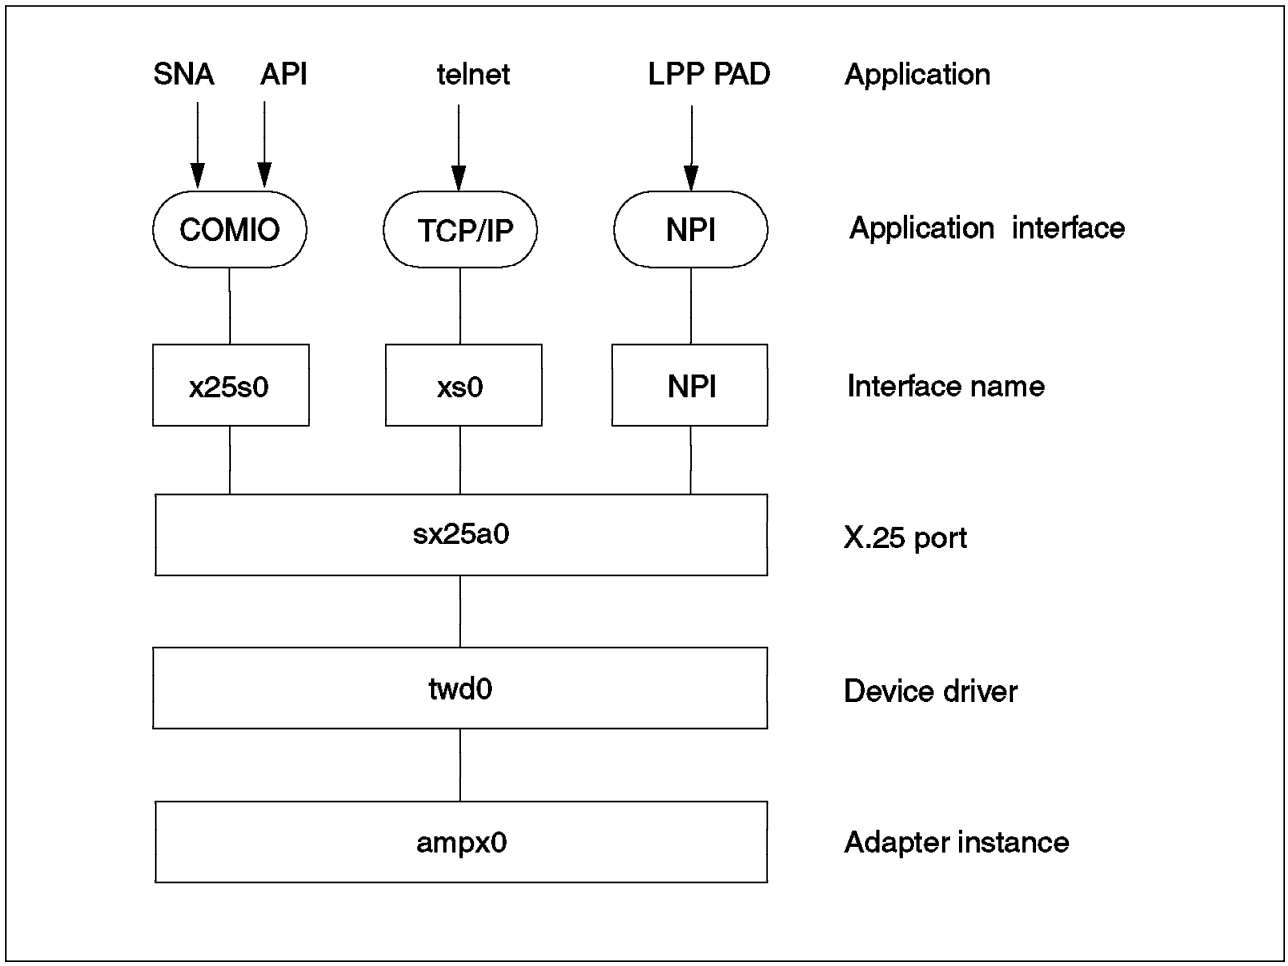

Figure 9. X.25 Drivers and Interfaces

COMIO Common Input/Output (COMIO) is an application interface that enables different applications, such as SNA or any previously used API program, to use the new X.25 product by providing an emulation of the old interface x25s0.

TCP/IP TCP/IP is now using the streams-based interface xs0 which allows faster communication and higher data transfer rates.

NPI Network Provider Interface (NPI) provides a programming interface for the packet layer. There are no separate configuration steps for NPI.

LPP PAD The new product provides an X.25 Packet Assembler-Disassembler (PAD) function as part of the product.

5.2.2 AIXlink/X.25 Packaging

As stated above, the new X.25 Licensed Program Product for AIX Version 4.1 is divided into four filesets, and these are divided into separately installable filesets. The names and descriptions of those filesets are listed in the tables below.

- AIXlink/X.25 Package

  This package provides the base X.25 function, including the protocol stack, device drivers, and adapter microcode. It also includes support for TCP/IP, NPI, DLPI, SNMP, and a PAD. A maximum of 512 virtual circuits for each line is supported.

Table 15. AIXlink/X.25 Package Contents

Fileset Name	Description
sx25.rte	AIXlink/X.25 Run-Time Environment This fileset provides the base X.25 device drivers, configuration methods, and applications necessary to use the other AIXlink/X.25 options.
sx25.npi	AIXlink/X.25 NPI and DLPI Support This fileset provides the device driver necessary to run Network Provider Interface (NPI) applications.
sx25.comio	AIXlink/X.25 COMIO Compatibility Support and Applications This fileset provides the device driver and applications for the COMIO compatibility interface. This interface provides compatibility with the AIX Version 3 base X.25 product.
sx25.tcpip	AIXlink/X.25 TCP/IP Support This fileset provides support for the TCP/IP protocol running over X.25.
sx25.pad	AIXlink/X.25 Triple-X (PAD) This fileset provides PAD software supporting the X.3, X.28 and X.29 standards.
sx25.server	AIXlink/X.25 Server Support This fileset is part of the standard AIXlink/X.25 package and provides an increased number of available virtual circuits based on the customer licensing agreement.

- AIXlink/X.25 Low Capacity Package

  This package provides the same functionality as the AIXlink/X.25 package, but with a reduced number of available virtual circuits (4). It is intended to be a low cost product for client systems.

  It contains the same filesets as AIXlink/X.25 above, except for the sx25.server fileset, which is not supported.

- AIXlink/X.25 Application Development Toolkit Package

  This package provides libraries, include files and sample programs that can be used in developing NPI, DLPI or COMIO applications for X.25.

Table 16. AIXlink/X.25 Application Development Toolkit Package Contents	
Fileset Name	**Description**
sx25.adt.rte	AIXlink/X.25 Application Development Toolkit - Run-Time Environment This fileset provides include files, libraries and sample programs for base X.25 program development.
sx25.adt.npi	AIXlink/X.25 Application Development Toolkit - NPI/DLPI This fileset provides include files, libraries and sample programs for NPI and DLPI program development.
sx25.adt.comio	AIXlink/X.25 Application Development Toolkit - COMIO Support This fileset provides include files, libraries and sample programs for COMIO compatibility program development.

- AIXlink/X.25 InfoExplorer Package

  The AIXlink/X.25 InfoExplorer package contains AIXlink/X.25 documentation for system administrators, application developers, and end users in the form of a hypertext information base.

Table 17. AIXlink/X.25 InfoExplorer Package Contents	
Fileset Name	**Description**
sx25.info.en_US.usr_gd	X.25 User's Guide This fileset contains information on using, managing, and programing the AIXlink/X.25 Version 1.1 application for AIX Version 3.2.5 and AIX Version 4.1.

5.2.3 X.25 General Migration Path

The general flow of an upgrade from X.25 running on AIX Version 3.2 to AIX Version 4 is slightly different from other upgrades. The steps are as follows:

1. Write down, print or save the X.25 adapter definitions
2. Write down, print or save the TCP/IP definitions
3. Export your SNA definitions
4. Upgrade the AIX system
5. Install the new X.25 product
6. Install the new SNA version
7. Re-enter, check or migrate the X.25 adapter definitions
8. Add the additional X.25 interface for COMIO
9. Add the additional X.25 interface for TCP/IP
10. Test the connection to the X.25 network
11. Migrate and test SNA
12. Migrate and test your API programs
13. Migrate and test shell scripts

The new steps are described in more detail below:

5.2.4 Saving Existing X.25 Adapter Definitions

The adapter definitions are stored within the AIX ODM data base. Depending on your installation method, they may disappear after the AIX upgrade. Therefore, it is always wise to save your old definitions. The method for saving depends on the X.25 support you are using on your AIX Version 3.2 system.

BOS X.25 Support There is no automatic saving function available. Document the old definitions by printing the adapter definition screens. One method for doing this could be:

1. Erase or rename the smit.log file.
2. Enter each X.25 configuration SMIT panel.
3. In each panel press F8, then Enter to save a copy of the screen into the smit.log file.
4. After saving each screen, print or save the smit.log file.

X.25 LPP Use the command:

```
backupx25 -d directory_name
```

in order to save all your X.25 adapter definitions in the specified directory. For later use, save these files on a diskette or tape.

5.2.5 Saving X.25 TCP/IP Definitions

The TCP/IP definitions are also stored within the AIX ODM data base. Again, depending on your installation method, you may need to recover these definitions after an AIX upgrade. Save the information by examining the following SMIT panels.

- X.25 Network Interface definitions
- X.25 IP Host definitions (translation table)

5.2.6 Saving X.25 SNA Definitions

The SNA definitions are stored within the SNA ODM data base and can be saved by the

```
exportsna
```

command. This is already described in 3.2.2, “Backing Up SNA Configuration Profiles” on page 75.

5.2.7 Installing New X.25 Product

After installation and migration of AIX Version 4.1, you have to install the new AIXlink/X.25 for AIX Version 4.1 product. AIXlink/X.25 is installed in the same manner as any other AIX LPP.

The minimal installation is simply:

sx25.rte AIXlink/X.25 Runtime Environment

The other packages to be installed will depend upon your configuration and license.

5.2.8 Restoring the X.25 Adapter Definitions

In order to get the new product into production, the definitions for the adapter and any applications using X.25 must be restored.

Depending on your old product, you either have to re-enter all of the definitions, or they can be restored from a previously saved file.

BOS X.25 Support In this case, no restore is function available. You must enter all definitions manually using the documentation of the old definitions printed in 5.2.4, “Saving Existing X.25 Adapter Definitions” on page 122. Refer to *AIXlink/X.25 1.1 for AIX: Guide and Reference*, SC23-2520, for details on configuring the new X.25 software.

Note: Since many parameters have changed in the new product, you should carefully read about these changes in the X.25 manuals.

X.25 LPP Use the command

```
restorex25 -d directory_name
```

to restore all your X.25 adapter definitions from the copy you previously saved in 5.2.4, “Saving Existing X.25 Adapter Definitions” on page 122.

Using the command:

```
lsdev -C
```

you should see the following devices available:

```
apmx0      Available 00-04          X.25 CoProcessor/2 Adapter
twd0       Available 00-04-00       N/A
sx25a0     Available 00-04-00-00    N/A
```

The command:

```
lsx25
```

will give you a report on the configuration of the X.25 ports. For example:

```
*****************************************************************
* Configuration report for X.25 LPP ports configured           *
*****************************************************************
Machine: isar12.ak.munich.ibm.com

*****************************************************************
* Report by slot number - bus 0
*****************************************************************
Slot 1, scsi0           SCSI I/O Controller
Slot 2, tok0            Token-Ring High-Performance Adapter
Slot 3, gda0            Color Graphics Display Adapter
Slot 4, apm0            6-Port Portmaster Adapter/A X.21
Slot 4, twd0            X.25 Streams driver
        Physical port 0 is x25 port sx25a0 [45890060292]
*****************************************************************
* Report by X.25 port's logical location                       *
*****************************************************************
X.25                                    Logical Logical
Port    Driver  NUA     COMIO   TCP/IP  board   port
sx25a0  twd0    45890060292     x25s0   n/a     0       0

*****************************************************************
* Report by X.25 port's physical location                      *
*****************************************************************
X.25                            Phys.
Port    Driver  Adapter Slot    Port    Interface
sx25a0  twd0    apm0    4       0       X.21
*****************************************************************
* Report by NUA                                                *
*****************************************************************
 NUA                    X.25 Port
45890060292             sx25a0
```

This report will continue with information about any configured COMIO and TCP/IP interfaces. The output depends on your configuration. In the sample above, a Portmaster/A adapter was used, which is supported by the new X.25 LPP only.

Note

The restorex25 process will restore all X.25 definitions, including COMIO and TCP/IP definitions. In this case, the next two steps can be skipped.

5.2.9 Adding the COMIO Interface

In order to use SNA or any API programs over X.25, you must have installed the sx25.comio fileset. To use it, the COMIO interface must be added through the SMIT fastpath `smit cx25str_mp`

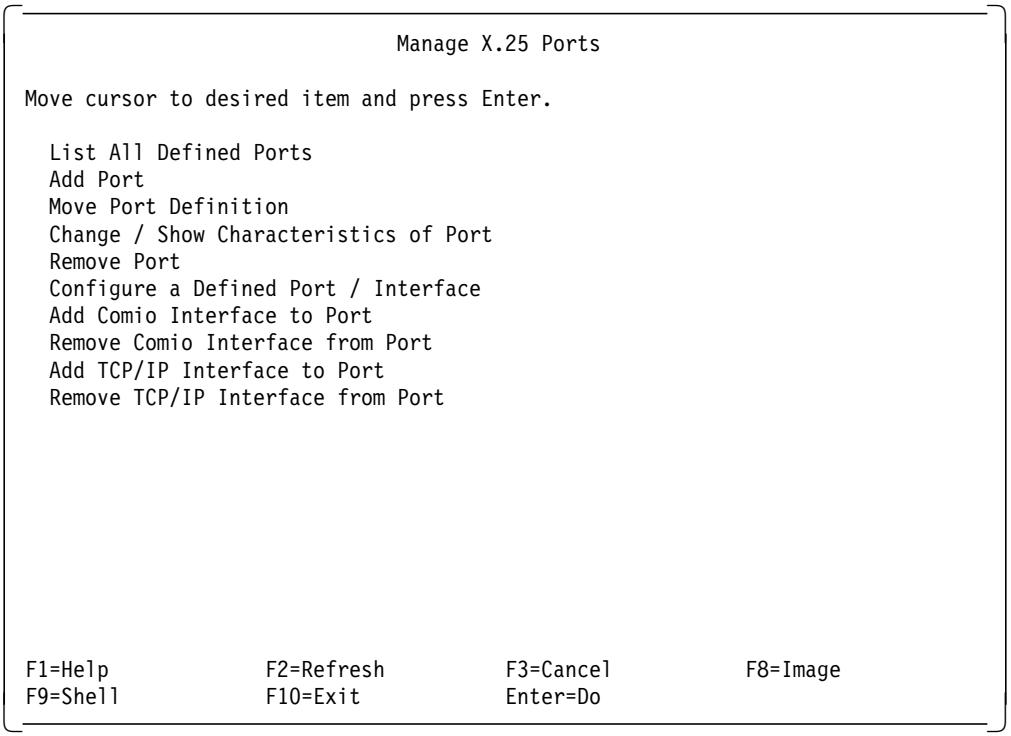

Select Add Comio Interface to Port.

5.2.10 Adding the TCP/IP Interface

In order to use TCP/IP over X.25, you must have installed the sx25.tcpip fileset. To use it, TCP/IP must be added by choosing the SMIT option **Add TCP/IP Interface to Port**. Use your documentation of the old definitions to enter the IP Address and network mask.

To get a mapping between the IP Address and the X.25 Network User (NUA) Address, an *IP Host Configuration* entry must be made using the:

```
smit mksx25
```

command which results in the following panel:

```
                          Add a X.25 LPP IP Host Entry

Move cursor to desired item and press Enter.

  Add a Permanent Virtual Circuit (PVC) X.25 LPP IP Host Entry
  Add a Switched Virtual Circuit (SVC) X.25 LPP IP Host Entry

F1=Help             F2=Refresh          F3=Cancel           F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Enter the values as you had them in your old X.25 product for each IP host you communicate with when using the X.25 network.

5.2.11 Testing the Connection to the X.25 Network

After adding the COMIO interface, you will be able to perform a basic network test. The simple application `xtalk` can be used to try connections either to a network-supplied test address or to any other NUA on the X.25 network.

If you have TCP/IP connections, you should also test these now.

5.2.12 Restoring the SNA Definitions

All SNA profiles must be migrated separately using the normal SNA migration path. From the X.25 standpoint, there are no changes, because COMIO provides the emulation of the old X.25 interface.

Do not forget to add the X.25 DLC fileset when installing SNA in order to get the *dlcqllc* device.

5.2.13 Migrating X.25 API Programs

The Common Input/Output (COMIO) interface emulates the device driver interface supported in previous releases of AIX X.25. This allows the user-space library API (available in AIX Version 3.2 BOS X.25 support) and any applications written to the device driver interface to work unchanged. Since the library is still available in the X.25 LPP, all applications written to use this library should work without recompilation.

You still need access to a C compiler, however, because the X.25 API includes a library of C subroutines that use the services of the COMIO emulator. These library routines must first be compiled. Application programs call these subroutines to access X.25 functions. The X.25 API subroutines are kept in the /usr/lib/libx25s.a library.

Note: Any new applications should be developed to use the NPI interface. NPI uses *streams* allowing faster access and faster data transfer for X.25 traffic.

5.2.14 Migrating Shell Scripts

In the X.25 LPP, many commands, attributes and definitions have changed from the AIX Version 3.2 BOS X.25 support. Therefore, any shell scripts referring to the old commands, such as `xmanage`, `xroute` and `xmonitor`, must be checked and altered as appropriate. For example, the IP host entry could be done by the command `x25xlate`, while now the `x25ip` command must be used. Also, many attributes for the X.25 definitions have changed. If your shell scripts refer to any attribute, it must be checked and changed. A complete description on all changes can be found in:

- *AIXlink/X.25 1.1 for AIX, Guide and Reference*, SC23-2520
- *AIX/6000 X.25 LPP Cookbook*, GG24-4475

Part 3. Migration of Multiple Systems

Chapter 6. Migrating Multiple Systems in a LAN Environment

The goal of this chapter is to outline concepts and procedures for migrating several systems at a time from AIX V3.2 to AIX V4.1 in a Local Area Network (LAN) environment. We present you with a migration method that will avoid any local interaction at the systems. No media handling or entering of commands at the system consoles will be necessary.

6.1 Migrating Multiple Systems Issues

The questions you have to look at when migrating multiple systems connected through a LAN are the following:

- Which AIX V4.1 installation method is the best for the migration?
- How do you back up all of the AIX V3.2 systems?
- How do you migrate AIX V3.2 installation servers?
- How do you reboot systems remotely in order to start an AIX V4.1 installation?
- How do you distribute the code that is needed during the migration (including BOS and additional LPPs) to the systems without the need for local media?
- How do you avoid the need for local manual interaction at the systems?
- Is there a tool that shows a reasonable performance doing all this?
- How do you migrate AIX V3.2 Diskless Workstation Management (DWM) servers?

In the following sections, we discuss these issues. First, we review briefly the different AIX V4.1 installation methods that can be used in a LAN environment with Network Installation Management (NIM) and choose the best one for our needs. Then we discuss how to back up the AIX V3.2 systems before starting the migration and how to migrate AIX V3.2 installation servers. After this, we describe the Network Installation Management (NIM) tool which is the suitable tool to fulfill our needs. In a special section, we discuss the migration of AIX V3.2 DWM servers. A scenario follows with detailed descriptions of all the steps which are necessary to prepare and run the migration of multiple systems connected through a LAN.

6.2 Choosing the Installation Method

The first thing you should do before starting a migration is figure out the method that best suits your needs for getting your systems to a new operating system level. We discuss briefly in this section the different installation methods you can use with AIX V4.1, and we point out what we think is the best way to migrate several AIX V3.2 systems on a LAN.

If you need more details about the different installation methods, please refer to *A Holistic Approach to AIX V4.1 Migration, Volume 1, SG24-4652.*

6.2.1 New and Complete Overwrite Install

The New and Complete Overwrite Install is the preferred installation method to install new systems from scratch. But it can also be used to reinstall existing systems if you want to wipe out everything on the disks you choose as destination disks for the installation. If disks not chosen for installation contain volume groups other than rootvg, they are not destroyed and can be imported later.

Attention!

Please remember: Your complete rootvg will be destroyed including all configuration information during a New and Complete Overwrite Install.

Therefore, we do not recommend the New and Complete Overwrite Install to migrate multiple AIX V3.2 systems on a LAN.

You may choose the New and Complete Overwrite Install if you want to wipe out everything in the rootvg of your systems and start from scratch. But we do not see this as a migration in the true sense of the word, and therefore you will not find detailed information about it in this chapter. However, you will find information on how to set up the Network Installation Management (NIM) for a New and Complete Overwrite Install in the section 6.10, “Setting Up NIM for New Installation” on page 157.

6.2.2 Preservation Install

The Preservation Install is the preferred installation method to reinstall existing AIX V3.1 or AIX V4.1 systems without destroying the user data. It recreates the four file systems, /, /usr, /var, and /tmp, so that all data stored there is lost, including the configuration files. Only /etc/filesystems is preserved so that you will be able to access your file systems in the rootvg after the Preservation Install (for example /home). Volume groups other than the rootvg are not destroyed either.

Attention!

Please remember that all data and all configuration information in the file systems, /, /usr, /var, and /tmp, except /etc/filesystems, are lost during the Preservation Install.

Therefore, we do not recommend the Preservation Install for migrating multiple AIX V3.2 systems on a LAN. You may choose the Preservation Install if you want to wipe out all configuration and system data for some reason but keep the user data. But be aware that you will have to configure the system from scratch again.

We do not see this as a migration in the true sense of the word, and therefore you will not find detailed information about it in this chapter. However, you will find information on how to set up the Network Installation Management (NIM) for a Preservation Install in section 6.10, “Setting Up NIM for New Installation” on page 157.

6.2.3 Cloning With mksysb Install

The term cloning describes the installation of a system backup (created with the mksysb command on one system) on different systems in order to make the software configuration of all systems as identical as possible.

However, with AIX V3.2, this has not always been easy mainly because of differences in the hardware between the systems. For AIX V4.1, it became even more difficult to ensure a proper cloning for several reasons. The two most outstanding reasons are that not all device drivers are installed automatically with AIX V4.1 and SMP systems require a different AIX kernel to run.

Regarding migration of multiple systems on a LAN, cloning the systems by installing a mksysb image is not the best choice. The paramount reason for this is the fact that during cloning, all existing system configuration data is wiped out and replaced with the configuration data of the mksysb image. So, it is basically a New and Complete Overwrite Install using another system′s image.

Attempts are made during the cloning to adjust the system′s hardware configuration in the Object Data Manager (ODM) with the one from the mksysb image, but this is not always possible. Another severe disadvantage is the fact that cloning between different systems is still not an officially supported installation method. The mksysb command was originally designed to get a backup from one system and to reinstall it on the same system, if needed. Nevertheless, cloning works in most of the standard cases. But if problems occur during cloning, there may be no support for it.

For our migration plans in a LAN environment, we do not see cloning as the preferred method to accomplish what we need. The main reason is that all configuration data is wiped out during cloning. To avoid this, you would need for every system a specific mksysb image containing the customized configuration of the system or a specific method to reconfigure each system. This would make the use of cloning rather complicated. Therefore, you will not find detailed information about it in this chapter. However, you will find information on how to set up the Network Installation Management (NIM) for cloning in section 6.11, “Setting Up NIM for Cloning” on page 158.

6.2.4 Migration Install

The Migration Install is the preferred installation method for migrating existing AIX V3.2 systems to AIX V4.1 while keeping the configuration and user data. It is an installation method that works on the basis of files. It does not recreate file systems, but works within the existing file systems and replaces system files there on an individual basis. Files added by the users are left untouched. Configuration files are migrated from AIX V3.2 to AIX V4.1 with different methods.

The Migration Install basically works in two steps. The first step is to reboot the system, and install the Base Operating System (BOS) and LPPs that come with BOS, such as X-Windows and TCP/IP. The second step is the installation of additional LPPs that must be purchased separately. Both steps migrate their components, including the configuration files, from an AIX V3.2 to an AIX V4.1 level during the installation process, usually without the need for additional manual intervention except in special cases.

For more details on the Migration Install method, please refer to *A Holistic Approach to AIX V4.1 Migration, Volume 1, SG24-4652.*

Looking at our goal to migrate simultaneously several AIX V3.2 systems to AIX V4.1 on a LAN without manual intervention, the Migration Install method seems like the right choice. It is able to migrate systems from AIX V3.2 to AIX V4.1 for BOS and other components while preserving the configuration and user data. We therefore chose it for the migration procedure outlined in this chapter.

6.3 Backing Up AIX V3.2 Systems

Please remember to first back up all systems you want to migrate so that you will be able to recover them as quickly as possible if something goes wrong during the process.

Operating a network of clients and servers, you probably already have a backup policy in place. In case you are looking for a method to back up all your clients and to recover them quickly, we recommend that you use the AIX V3.2 Network Installation Server mechanism to store and distribute your AIX V3.2 `mksysb` images. Network Install Management (NIM) cannot handle AIX V3.2 `mksysb` images.

You can back up your clients to file systems mounted over NFS to servers. Later, you can use these servers to re-install the `mksysb` images on the clients, if necessary. This is even possible with servers already running AIX V4.1 and NIM. The AIX V3.2 Network Installation Server mechanism can be installed to AIX V4.1 systems with the fileset bos.compat.NetInstl. It can be run in parallel with NIM; in fact, it is a totally separate functionality.

There is a good description of the setup in InfoExplorer. You can find it if you search for ″network installation server″.

Attention!

Remember that you need to create boot diskettes for the AIX V3.2 clients as described in the AIX V3.2 file /usr/lpp/bos/README, or have another bootable medium available, such as a `mksysb` tape. The AIX V3.2 Network Installation Server is not capable of booting systems over the network the way NIM can do it.

In case you have AIX V3.2 installation servers that you want to migrate to AIX V4.1, please see section 6.4, “Migrating AIX V3.2 Installation Servers” for details.

6.4 Migrating AIX V3.2 Installation Servers

If you are using one or more AIX V3.2 installation servers for your LAN environment, you may want to migrate these servers to AIX V4.1 to be able to use them as AIX V4.1 installation servers for the migration of the other systems on the LAN. This should be the first step before migrating all other systems.

Attention!

It is possible to migrate AIX V3.2 installation servers to AIX V4.1 installation servers and still be able to serve AIX V3.2 software and `mksysb` images to clients if needed.

This section describes what to consider during migrations of AIX V3.2 installation servers.

When you set up an AIX V3.2 installation server, you basically do the following:

1. Configure TCP/IP and NFS.
2. Create a netinst user.
3. Configure the `inetd` subserver instsrv.
4. Create an extra file system and fill it with AIX V3.2 installation images.
5. Create the choices file /home/netinst/db/choices.
6. Create class or client description files (optional).
7. Export the installation image file system to the clients with NFS.

In AIX V4.1, the functionality to serve AIX V3.2 installation images is possible by installing the compatibility fileset bos.compat.NetInstl. AIX V4.1 installation images are served by NIM (see section 6.5, “What is NIM?,” for details on NIM). It is possible to use NIM and the AIX V3.2 installation server mechanism in parallel to serve both AIX V4.1 and AIX V3.2 software from an AIX V4.1 system.

To migrate an AIX V3.2 installation server to AIX V4.1, you need to use the Migration Install method. During the Migration Install of a system serving AIX V3.2 installation images, the installation process senses that the system is an AIX V3.2 installation server and automatically installs the bos.compat.NetInstl fileset to keep the AIX V3.2 installation server functionality. During the migration, all the configuration and data files you created in AIX V3.2 with the seven steps above are preserved except the choices file, /home/netinst/db/choices, from step 5. It is overwritten by a default file.

To recreate the choices file, use the command:

```
# echo '<installation images directory>/*' > /home/netinst/db/choices
```

After this step, you are able to serve AIX V3.2 installation images from an AIX V4.1 system as usual. The AIX V3.2 installation server is migrated to AIX V4.1 and is still able to serve AIX V3.2 installation images in addition to the AIX V4.1 installation images that can be served by NIM.

That means that if you store AIX V3.2 `mksysb` images on an AIX V3.2 installation server, you will be able to migrate the server to an AIX V4.1 installation server, and it will still have your AIX V3.2 `mksysb` images stored and will be able to serve them to clients.

6.5 What is NIM?

Network Installation Management (NIM) is a tool that comes bundled with AIX V4.1 at no additional cost. It is a replacement for the AIX V3.2 Diskless Workstation Management, but can do more than this. The working environment of NIM are LANs. There is one master that controls the NIM environment, the servers that hold the necessary resources and the clients that are being served. In this section, we look at NIM in general, show the features, the structure, the basic setup, and the limitations.

6.5.1 NIM Features

NIM is able to:

- Initialize and serve diskless and dataless clients including remote booting
- Install BOS for stand-alone clients, including remote booting
- Install `mksysb` images for stand-alone clients, including remote booting
- Install other software products, like LPPs and PTFs or even maintenance levels to clients (this feature is new in AIX V4.1.3); bundling may be used during installation
- Perform software maintenance for clients; for example de-installing of software or checking for PTFs (this feature is new in AIX V4.1.3)
- Start scripts as part of the installation process on clients
- Use special setup files to achieve unattended installation on clients

6.5.2 NIM Structure

NIM works with three different object classes to accomplish all the tasks.

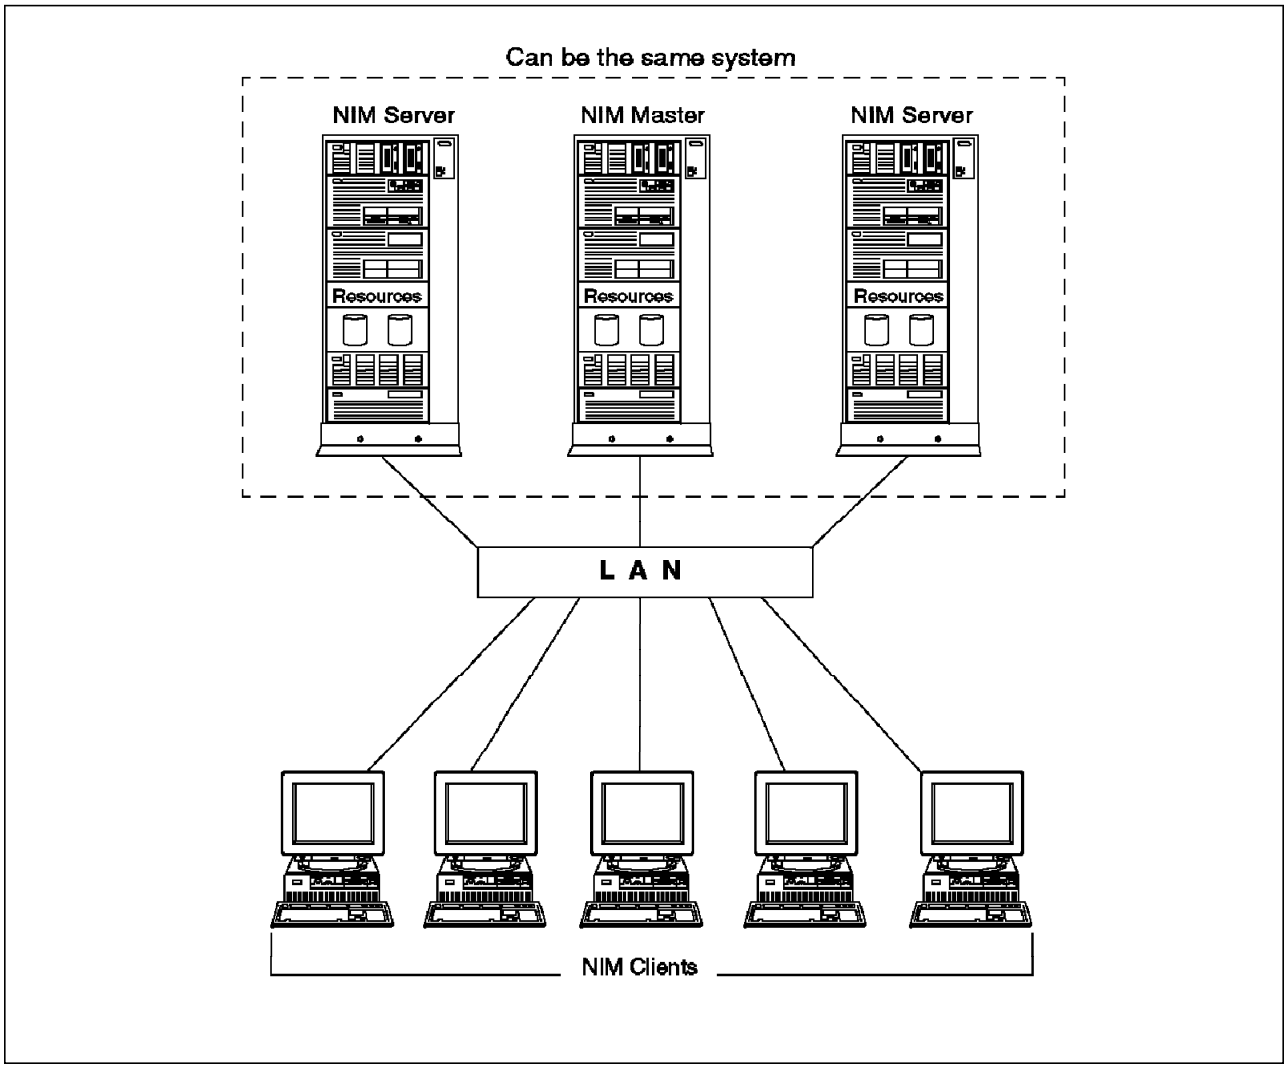

Figure 10. NIM Structure

6.5.2.1 Network Objects

Network objects are used to define the LAN subnets in a NIM environment. NIM needs a network object for every subnet that will participate in the NIM environment. NIM also needs its own routing definitions in addition to any TCP/IP routing that may exist. Routing is defined within NIM by giving NIM the two TCP/IP host names of the gateway that subnets use to communicate with one another. NIM uses this information to establish TCP/IP routes during AIX installation. Before using NIM, the TCP/IP network must already be configured. At minimum, you have one network object to which the NIM master is connected. This network object is automatically created when you configure the NIM master.

6.5.2.2 Machine Objects

Machine objects are used to define the systems that take part in the NIM environment. The master machine object is automatically created when the master is configured. There can only be one master in a NIM environment. All servers that serve resources to clients and all clients must also be defined as machine objects. Clients can be of the type stand-alone, diskless or dataless. Servers must be of the type stand-alone. The NIM master can also serve resources at the same time.

6.5.2.3 Resource Objects

Resource objects define the resources needed by the NIM master in order to perform the installation operations on the clients (machine objects) in the NIM environment. Operations include BOS installation, diskless/dataless initialization, software maintenance and software customization (LPP distribution). There are different types of resources: spot, installp_bundle, mksysb, script, dump, paging, boot and other. The use of these resources depends on the kind of operation you are going to perform and on the machine types (stand-alone, diskless, dataless).

Resources must be created and allocated to the clients before commencing installation operations.

6.5.2.4 Operations

The following are some of the more significant NIM operations.

- bos_inst

  Performs a BOS installation or migration, including reboot

- cust

  Performs software customization, including installation of additional filesets and filesets updates

- maint

  Performs software maintenance including de-installation of filesets and commit or reject of fileset updates

- reset

  Resets the states of clients in case of problems

- check

  Checks the status of NIM machines and certain resources (spot and lpp_source)

- dkls_init

Initializes a diskless machines′s file systems

- dtls_init

  Initializes a dataless machines′s file systems

There are some more operations which may be needed in special cases.

6.5.2.5 Examples of Required Resources

Different client types (diskless, dataless, stand-alone) require different kinds of operations (dkls_init, dtls_init, bos_inst) to be initialized/installed. The required resources depend on the operation and the client type. Two examples follow for required resources:

1. BOS installation for a stand-alone client:

   Before starting the bos_inst operation on a stand-alone client, the following resources must be created and allocated to the stand-alone client:

   - spot

     Provides network boot support for migration via NIM.

   - lpp_source

     This resource contains the software packages used to install a stand-alone client during migration.

   - There are other optional resources. For example, the bosinst_data resource provides a special setup file enabling unattended installation of clients.

     Note: This has been introduced with AIX V4.1.3: if you use the force_push attribute with the bos_inst operation to remotely install a client running AIX V3.2 without local intervention at the client, a bosinst_data resource must be allocated to the client in any case.

2. Initialization of a diskless client:

   Before invoking the dkls_init operation, the following resources must be created and allocated to the diskless client:

   - spot

     This resource provides the /usr file system for the diskless client on a remote server.

   - root

     This resource provides the / file system for the diskless client on a remote server.

   - dump

     This resource provides the dump device for the diskless client on a remote server.

   - paging

     This resource provides the paging device for the diskless client on a remote server.

   - There are other optional resources. For example the home resource provides the /home directory for the diskless client. If it is not allocated, the /home directory will be created in the / file systems of the client and will not be a separate file system.

6.5.3 Setting Up NIM: Basic Workflow

The basic workflow of a NIM environment setup is shown in this section.

> **Attention!**
>
> We recommend that you use extra file systems or maybe even volume groups for some resources, like the lpp_source, as they may get quite big (more than 300 MB). Later, that makes the administration a lot easier.

1. Install the NIM filesets

   On the NIM master, the following filesets must be installed:

   - bos.sysmgt.nim.master
   - bos.sysmgt.nim.client (installed automatically as a prerequisite)
   - bos.sysmgt.nim.spot

     NIM′s spot creation process automatically installs bos.sysmgt.nim.spot on a spot server.

   TCP/IP and NFS are installed automatically as a prerequisite. Before you set up your NIM environment, TCP/IP and NFS must be configured and running. This includes definitions of networks, initialization of gateways, name serving, and routing. All the host names for the systems participating in the NIM environment (master, servers and clients) must be resolvable.

2. Configure the master fileset

   On the system which is planned to be the NIM master, you must configure the master fileset. The network object the NIM master is connected to is automatically created. All the following definitions are done on the NIM master.

3. Define additional networks and routing

   If you have several subnets, corresponding NIM networks must be defined, including the routing between them.

4. Define clients and servers (via machine objects)

   Each system participating in the NIM environment as a client or a server must be defined on the NIM master. Of course, you can add systems later when your environment grows.

   NIM servers other than the NIM master must be installed and configured as stand-alone NIM clients. That means that they must run AIX V4.1, and the corresponding TCP/IP, NFS and NIM filesets must be installed and configured.

   You can find information on how to set up such NIM servers in section 6.12, “Migration Scenario” on page 159.

5. Define resources

   All necessary resources for the operations must be defined. Remember that resources can be distributed on machines other than the master in the NIM environment.

6. Allocate resources

The resources for the planned operations must be allocated to the clients. What resources must be allocated depends on the operations you want to start and on the client type.

7. Start operations on clients

   Now the operations on the clients can be initiated (for example a BOS installation operation bos_inst).

   After AIX V4.1 BOS is installed on a system via NIM, the system will be, by default, automatically set up as a NIM client and will be able to receive commands from the NIM master. That means that TCP/IP and the bos.sysmgt.nim.client fileset are automatically installed and configured via NIM.

Although this may look a bit complicated, setting up NIM using SMIT is quite easy when the basic flow is understood. We describe the flow for a Migration Install with NIM in section 6.9, “Setting Up NIM for Migration: Flowchart” on page 155.

In section 6.12, “Migration Scenario” on page 159, we describe in detail how we set up and used NIM in our LAN scenario.

6.5.4 NIM Limitations

Even though NIM is capable of performing a lot of operations on clients on a LAN, it has some limitations that are important to know. We describe them here in detail because you may consider using other tools, such as Software Distribution for AIX. This list should give you a basis to make a comparison with other tools:

- NIM only supports the LAN network types token-ring, Ethernet and FDDI. No other network types are supported.
- NIM only supports TCP/IP and NFS; no other protocols, like AFS or SNA, can be used.
- The network definition in NIM, including routing, must be done manually. NIM is not able to automatically adopt the TCP/IP network and routing definitions.
- NIM masters and servers must be systems running under AIX V4.1.
- NIM can only distribute AIX V4.1 BOS and `mksysb` images. Under AIX V4.1, there is still the possibility to set up the AIX V3.2 network installation mechanism by installing a compatibility fileset. This way, you can distribute AIX V3.2 `mksysb` and BOS images from an AIX V4.1 server. But it is a functionality totally outside of NIM that has nothing to do with it. However, AIX V3.2 LPPs can be distributed and installed with NIM.
- NIM can only distribute and install code that uses `installp` for installation, that means it is limited to AIX and cannot serve other architectures. The advantage is that it is specialized for AIX and can do everything that `installp` can do.
- Each system participating in the NIM environment as a server or as a client must be defined one by one. For large environments, this results in a lot of definition work to be done.
- NIM does not have the possibility to group systems, such as starting a BOS installation for several clients at a time. This can only be done by putting the high-level commands used by NIM in a shell script.

- There is no scheduling mechanism inside NIM (for example, to schedule an installation for a later time). This can only be done by scheduling the high-level commands NIM uses with the AIX mechanisms (like, for example, the `at` command).
- NIM cannot boot 40P systems without manual intervention. However, NIM can boot other PCI-based systems, such as the 43P, since the introduction of AIX V4.1.4.
- NIM cannot be a /usr server for clients; that means that the installation type /usr server is no longer known with AIX V4.1.
- It is not easy and not very well documented to change the IP address or the subnet mask of the NIM master once it is configured. The reason is that NIM adopts these values to use them in its own network description. One way to get it changed is to reinstall NIM and to reconfigure it.

6.6 Why NIM?

In this section, we check our requirements for migrating multiple AIX V3.2 systems on a LAN without manual intervention against NIM′s capabilities.

6.6.1 Booting AIX V3.2 Clients Remotely

We need a remote reboot capability to start the migration of AIX V3.2 BOS on the clients. Fortunately, NIM introduced a new feature called force_push with AIX V4.1.3. You can read the details about it in the file /usr/lpp/bos.sysmgt/nim/README on systems running AIX V4.1.3 and later that have the bos.sysmgt.nim.master fileset installed.

Basically, force_push allows you to ″push″ the reboot and installation of AIX V4.1 to client systems running AIX V3.2 from the master without local manual intervention at the clients. This even works with older systems that do not have an IPL ROM capable of booting from network devices (for example the IBM RISC System/6000 model 320, 520 and 530).

Three prerequisites must be met in order to use the force_push mechanism:

1. The AIX V3.2 clients must have an .rhosts file in the / directory granting root access to the NIM master.
2. Key must be in Normal position on the target system.
3. The NIM master must use the force_push attribute during bos_inst operations.

6.6.2 Distributing AIX V4.1 Code to Clients

We need a code distribution mechanism that is capable of supplying the clients with all the software they need for BOS and LPP migration without the need for any local media.

NIM can distribute and initiate the installation of the AIX V4.1 BOS on the AIX V3.2 clients. Once the AIX V4.1 BOS is installed, LPPs, PTFs and maintenance levels can be distributed. This meets our requirements completely.

6.6.3 Starting the Migration without Local Client Interactions

To achieve an unattended migration of the clients, we need a way to initiate a BOS migration and additional LPP migrations from AIX V3.2 to AIX V4.1 without any manual interventions or questions at the clients.

With the use of the bosinst_data resource to define the installation parameters for BOS, it is possible to start a BOS migration without manual intervention at the clients. With the installp_bundle resource, we can define what additional LPPs or PTFs will be installed after the BOS migration. This all runs in one step without the need for additional interaction.

6.6.4 Performance and Sizing

Looking at performance, you have to consider two main points in a NIM environment:

1. The NIM servers
2. The network

NIM is able to distribute resources to several servers and to control and use several servers at the same time to achieve better performance during client installations.

The lpp_source resource and the spot resource are the most critical resources affecting performance during migrations. For details on resource types and how to set up NIM servers other than the NIM master, please refer to section 6.12.4.5, “Setting Up Other NIM Servers” on page 167.

Since NIM is working with NFS to serve resources, you should size your environment from an NFS point of view. First, do some simple calculations to determine if you have enough network bandwidth for the desired installs.

You first need to set up a NIM server as you will be using it for the final installs. Then you need to install a single client, and measure the total network data that is transferred between the NFS client and server during the installation. It is best to use a network analyzer for network tracing. If something like that is not available, then the `iptrace` command can be used, but it influences the server and needs some set up to report the total data transferred.

Then the total data transferred can be used to calculate the network requirements to install the systems. If you only have one server doing all of the installs and it has only one network interface, then it is obvious that it will be a bottleneck. Based on the total data transferred, you can calculate how much of a potential bottleneck it is going to be and how many servers and what LAN capacity you will need.

The other factors that you have to consider are server memory and disk layout. The more memory you have on the NFS server the better. This will increase the ability to cache the files that are being read from the server during the install process. Since the clients will be installing the same or similar software, the NFS server will have a better chance at having the data in memory and will be able to respond more quickly to the clients. This will also increase the throughput of the NFS server based on the fact that potential disk operations can be avoided.

Unfortunately, it is not possible for us to give you specific numbers on NIM servers and network sizing. Too many parameters are involved to make theoretical predictions.

6.6.5 Conclusion

NIM is an appropriate tool to perform an unattended migration of multiple AIX V3.2 clients to AIX V4.1 connected through a LAN. We therefore proceed and give you detailed information on how to set up NIM for migration and show you how we used NIM in our scenario.

6.7 Migrating AIX V3.2 DWM Servers

If in your installation you have diskless and/or dataless workstations, you will find information in this section on how to migrate the AIX V3.2 Diskless Workstation Management (DWM) servers.

In AIX V3.2, the tool used to set up and manage diskless or dataless workstations is called DWM. To set up an AIX V3.2 DWM server, the following steps must be run:

1. Initialize TCP/IP and NFS.
2. Start the `bootpd` and `tftpd` daemons.
3. Initialize the SPOT (Shared Product Object Tree) and install software in it.
4. Configure the diskless clients.
5. Convert the diskless clients to dataless clients, if necessary.

With AIX V4.1, NIM is used to serve diskless or dataless clients. Please refer to section 6.5, “What is NIM?” on page 135, for details on NIM.

Attention!

There is no defined migration path for migrating an AIX V3.2 DWM server to an AIX V4.1 NIM server. The DWM configuration information is lost and must be recreated in NIM.

During our tests, we used a Migration Install to migrate an AIX V3.2 DWM server to an AIX V4.1 NIM server. The server used two extra file systems, /tftpboot and /export, for the DWM resources.

Even though server configuration data, in addition to DWM configuration, is kept during the Migration Install, the DWM environment must be manually recreated in NIM.

We recommend that you run the following steps on the server to migrate it to AIX V4.1:

1. Do a system back up in order to be able to recover your system if something goes wrong.
2. Run a Migration Install on the server to get it to AIX V4.1 by booting from AIX product tape or CD-ROM and selecting the migration option from the BOS install menus.
3. Erase all NFS export definitions that have been used by AIX V3.2 DWM, for example with `smitty rmnfsexp`.

4. Erase the old boot images (typically in /tftpboot).
5. Erase the old dump, paging and share directories used by AIX V3.2 DWM (typically /export/dump, /export/swap and /export/share). No data that you may want to keep should be stored there. Keep the old /export/exec, /export/root and /export/home directories. User data that you may need later is stored there.
6. Install the NIM filesets bos.sysmgt.nim.master, bos.sysmgt.nim.client and bos.sysmgt.nim.spot.
7. Be sure that all NFS subsystems, including rpc.statd, are running on the server. You can check them with the command:

   ```
   # lssrc -g nfs
   ```

   If they are not running, you can start them with the command:

   ```
   # startsrc -g nfs
   ```

   You can modify the NFS startup file, /etc/rc.nfs, to start them after each reboot.
8. Configure the server as the NIM master.
9. Define the NIM diskless/dataless clients.
10. Define the spot, root and home resources in the directories /export/exec_new, /export/root_new and /export/home_new. /export/exec_new must be created first. /export/root_new and /export/home_new are created automatically. We cannot use the old directories for these resources because we either get errors, or they are erased during initialization.
11. Define the paging (only for diskless clients), dump and tmp resources in the directories /export/swap, /export/dump and /export/tmp. The directories are created automatically.
12. Allocate the resources, and run the NIM dkls_init or dtls_init operations to initialize the clients.

    Note: During the dkls_init/dtls_init operations, you need to make sure that your file system that holds the root and/or paging resources (typically /export) has enough space to hold the resources. It is not automatically expanded if more space is needed. Also note that paging size can be specified with the ″size″ attribute during dkls_init/dtls_init operations.
13. Boot the diskless/dataless clients.
14. Copy the client data that you want to keep from the directories /export/exec, /export/root and /export/home to the new directories /export/exec_new, /export/root_new and /export/home_new. When you copy data from /export/exec and /export/root, be sure that you do not copy AIX V3.2 system executables or system data, just user files or user configuration data (for example, /export/exec/<spot_name>/usr/local/bin). Be sure to copy the client′s home directories from /export/home/<client_name> to /export/home_new/<client_name>.
15. From the clients, test if all copied data is there.
16. Erase the old AIX V3.2 DWM directories /export/exec, /export/root and /export/home on the server.
17. Back up your AIX V4.1 NIM master.

Attention!

Seen from the diskless/dataless clients, this procedure is actually a new installation. No configuration data is preserved at the clients except the data copied between the directories. Therefore, all configuration information on the clients, for example extra file systems or user configurations, is lost and must be recreated. On the server, configuration data other than DWM data is preserved.

Since the majority of installations are stand-alone installations, we will not go into further detail concerning diskless/dataless server migration. The following sections all refer to migrations of stand-alone clients.

6.8 Setting Up NIM for Migration: General Considerations

This section contains the general points you have to consider when setting up NIM for a migration from AIX V3.2 to AIX V4.1.

6.8.1 Push Booting AIX V3.2 Clients

Since the AIX V3.2 systems do not have the NIM client fileset installed and configured and are not NIM clients yet, they cannot be reached via the standard NIM mechanisms. The only way to push-boot them is to set up an .rhosts file in their / directory granting root access to the NIM master, and use the force_push attribute during the NIM BOS installation operation on the NIM master.

Attention!

As stated in the file /usr/lpp/bos.sysmgt/nim/README, you need a bosinst_data resource in order to use the force_push attribute for the NIM BOS installation operation.

In our case, this is not a problem because we set up the bosinst_data resource anyway to achieve an unattended migration (see next point).

6.8.2 Setting Up bosinst_data Resource

The bosinst_data resource is a pointer to a modified bosinst.data file that defines parameters for the BOS installation, such as the console, the target disks, the locales, and so on. In this file, you also define the installation method you are going to use. In our case, we specified the following:

```
INSTALL_METHOD = migrate
```

You can find a template on AIX V4.1 systems under the path name /var/adm/ras/bosinst.data that you should copy to your NIM directories and then modify for your needs.

For unattended migration of the clients, you need to set up and allocate bosinst_data resources to each client. You will probably need several bosinst_data resources with different setups because not all clients have the same type of console (some will have graphic displays and some ASCII terminals), or you may need some other specific setup.

For a migration, the target disks do not need to be defined because the migration will automatically choose the disks where the rootvg is located.

Attention!

If your bosinst.data files miss important parameters or do not match with client configurations, the installation process reverts to prompting at the client consoles. You should watch the client consoles carefully during the first migrations to make sure that you do not miss any prompts.

For details on the bosinst_data setup we used, please refer to section 6.12, “Migration Scenario” on page 159.

6.8.3 Checking Ethernet Card Levels on the Clients

We found out during our migration tests that in order to be able to push-boot clients from the NIM master, the 3COM Ethernet cards in the clients must at least have a ROS level of 9. You should check this with the command `lscfg -v` before starting the migration. Check for the field called ROS Level and ID.

Attention!

If you have earlier levels, you will see a blinking 888 during the network boot on the client and get the codes 103-203-280 when pressing the reset button.

6.8.4 BOS Installation Operation Sources

Usually, you can use one of three sources for the NIM BOS installation (bos_inst) operation for a client:

1. The spot resource (source = spot):

   The spot resource serves three purposes in a NIM environment:

   a. It provides a /usr file system to diskless/dataless clients.

   b. The boot images for remote boots of the clients are generated from it.

   c. It acts as the installation source for a particular mode of BOS installation.

   When using the spot resource as the source for a BOS installation, all the files from the /usr directory in the spot resource are copied to the /usr file system on the client. That means that all the software that is installed in the spot resource is copied over to the client. The client might receive more software than necessary, resulting in an installation that is quick but big.

   On the server, you can either use the server′s /usr file system as a spot resource, or you can create an extra directory or an extra file system. The advantage of using the server′s /usr file system is that you save disk space on the server because files are not duplicated. The disadvantage is that you should no longer use the standard `installp` commands and interfaces (for example VSM) to install software into /usr. You have to use NIM instead, and that is often more complex. If you use an extra directory or file system, you do not have this restriction, but this will require more disk space.

   We recommend the use of an extra file system for the spot because that makes the administration of the servers a lot easier later.

2. The lpp_source resource (source = rte):

   The lpp_source resource is a directory or a file system where all the Backup File Format (BFF) images needed for the BOS installation of the clients are stored. If all required images exist, a special attribute, called simages, is set for the lpp_source. This way, you can check if the lpp_source is complete.

When rte is specified as the source parameter for the BOS installation, the BOS files are unpacked from the BFF images stored in the lpp_source resource. With this method, NIM can pick and choose the needed software (for example for device support), resulting in a smaller runtime BOS installation on the client. However, the installation takes longer than using the spot resource because the software must be unpacked.

3. The mksysb resource (source = mksysb):

   The mksysb resource is simply a mksysb image stored on a server that you can use for the installation of the clients.

   Using the mksysb resource, the same thing happens as if you used an mksysb image from a tape to install a client. All software in the image is put on the client, and the installation procedure tries to adjust the client's hardware configuration with the one in the image. So, it is not only a BOS installation but a reinstallation of a whole backup image onto the client.

Attention!

For a migration, only the lpp_source can be used as the source for the BOS installation.

But the spot resource must also be created and allocated to the clients before the migration can start. The reason is that the boot images that are generated from it are needed for the network boot of the clients.

6.8.5 Installing Additional LPPs

During a Migration Install started by NIM, only BOS and related LPPs (for example, TCP/IP or NFS) are migrated automatically on the clients.

For the migration of additional LPPs that do not belong to BOS (for example HCON or SNA), installp_bundle resources should be defined and allocated to the clients. They basically contain lists of all the additional LPPs that must be installed by NIM on the clients after the BOS migration. Usually, the lists are unique for every client, except if the amount of software on all clients is exactly the same. You should make sure that all the necessary software is available in the lpp_source resource before starting the migration (see next point).

In the installp_bundle resource, you can use lists of filesets or lists of package names of the extra LPPs. Both can be obtained from the table of contents (TOC) of your media. In section 6.12.4.13, "Creating the installp_bundle File" on page 181, you will find detailed steps on how to get this information.

Note: Only LPPs that have already been installed in AIX V3.2 are migrated during the Migration Install. LPPs listed in the installp_bundle files that were not installed in AIX V3.2 are just newly installed without any migration or customization work.

This means that it is possible to add new AIX LPPs during a NIM client migration. This can be achieved by simply listing these LPPs in the installp_bundle resources and by placing their images into the lpp_source resource.

6.8.6 Software in the lpp_source Resource

Another point that you should consider during the migration of AIX V3.2 clients with NIM is the fact that the lpp_source you choose as the source should contain all the software you need for migrating the clients, including extra LPPs that do not belong to BOS.

You will be able to run the whole migration of BOS and the LPPs in one step without any additional interaction using the installp_bundle resources only if you place them all into the lpp_source with the `bffcreate` command before starting the migration. Only then will the LPPs listed in the installp_bundle resources be automatically migrated with their requisites.

Attention!

Make sure that all the necessary requisites are loaded into the lpp_source before starting the bos_inst operation so that all LPPs on the clients can be migrated.

You can find out which requisites a fileset needs by looking into the TOC of your media (see section 6.12.4.13, “Creating the installp_bundle File” on page 181 on how to get the TOC) or by using the preview option of the `installp` command on an AIX V4.1 system.

Attention!

During the creation of an lpp_source resource, NIM only checks for the existence of a minimum set of support images that are required for client runtime installations. This set of images is not sufficient for migrations. For example, the C runtime, xlC.rte, is missing.

A list of the NIM support images is given in section 6.8.7.1, “Working with Product Tape” on page 149.

Because the dependencies between requisites can become quite complex, we recommend another approach. You can avoid requisite errors if you simply load all the software from your product tapes or CD-ROMs to the lpp_source. This way, you probably sacrifice some disk space on the servers, but you eliminate a potential source for errors during installation.

In case you suspect that you experienced missing filesets or requisite errors during the migrations, you can check the /var/adm/ras/devinst.log files on the clients. Here, you can find such errors under the sections:

`Missing Filesets`

or

`Requisite Failures`

6.8.7 Setting Up lpp_source Resource

You may set up a NIM master and NIM servers from a product tape or from a product CD-ROM. If using a CD-ROM, you need an AIX V4.1 Server CD-ROM to install the bos.sysmgt.nim.master fileset.

There is one major difference between the set up from tape and the one from CD-ROM. Working with a tape, you always need to copy BFF images required

for the lpp_source resource to disk. Working with a CD-ROM, you can define the lpp_source by just pointing to the mounted AIX V4.1 Server CD-ROM without copying the BFF images and save disk space. However, this method has some disadvantages that we describe in the next sections.

6.8.7.1 Working with Product Tape

Because NIM cannot access a tape directly as an lpp_source, you always need to copy the minimum set of support images required to fill an lpp_source from the tape to disk. You can do the copying during or before the creation of the lpp_source. If you use the tape as the source device for the install images during the lpp_source creation, all the required support images will be copied to disk during the creation. You can specify other BFF images that should also be copied to disk.

As clients usually have additional LPPs other than BOS that need to be migrated, we recommend that you do it the following way. First, copy all the required BFF images to a new directory or a file system on the disk using the `bffcreate` command; then define the lpp_source in this directory without giving a source for the install images. NIM only checks if the given directory contains all the necessary support images and defines the lpp_source pointing to that directory. This way, you can make sure that you have all the LPPs you need for your clients available in the lpp_source.

In case NIM is missing required support images in the directory you want to use as the lpp_source, you will get an error message during the definition of the lpp_source. The following images must be there:

- bos
- bos.net
- bos.rte.up
- bos.rte.mp
- bos.diag
- bos.sysmgt
- bos.terminfo
- bos.terminfo.data
- devices.base
- devices.buc
- devices.graphics
- devices.mca
- devices.scsi
- devices.sio
- devices.sys
- devices.tty

You can also check for the simages attribute with the command:

```
# lsnim -l <lpp_source_name>
```

If the line:

```
simages      = yes
```

is not present in the output, support images are missing.

6.8.7.2 Working with CD-ROM

When working with an AIX V4.1 Server CD-ROM, you can define the lpp_source pointing to the mounted CD-ROM by using the following parameters in the NIM SMIT menus or with the NIM commands:

- Specify the path name /<CD_mount_point>/usr/sys/inst.images as the location for the lpp_source.
- Do not specify a source for the install images to be copied from.

This way, NIM only checks if the necessary support images for the creation of the lpp_source are on the CD-ROM and uses the mounted CD-ROM as the lpp_source without copying BFF images to disk. This may save you more than 300 MB of disk space.

The disadvantage is that NIM only allows one lpp_source for the BOS installation of a client. If you use the mounted AIX V4.1 Server CD-ROM as the lpp_source, you need to make sure that all the software your clients need is on the CD-ROM. If some additional LPPs require a different CD-ROM, they will not be migrated or installed.

Attention!

Usually, clients have more LPPs installed than just BOS. Therefore, we do not recommend using the mounted AIX V4.1 Server CD-ROM as the lpp_source. Our goal of a one-step migration for all LPPs may not be fulfilled this way.

Rather, we recommend that you copy all the needed LPPs images from the AIX V4.1 Server CD-ROM and from other CD-ROMs if necessary with the `bffcreate` command to a new directory or file system the same way you would do it with a product tape. When creating the lpp_source, you can use this directory that already contains all the LPPs as the location. Without giving a source for the install images, NIM again checks only if all required support images are there and defines this directory as the lpp_source. Only by using this method can you make sure that all the necessary BFF images are available in the lpp_source.

The list of required support images for the lpp_source we gave you in the previous section 6.8.7.1, “Working with Product Tape” on page 149, and the methods to check for them also apply here.

This is basically the same method as using a tape. Therefore we will work with a common flow for tape and CD-ROM in section 6.9, “Setting Up NIM for Migration: Flowchart” on page 155.

6.8.7.3 Using CD-ROM to Save Disk Space

The default set of support images copied by NIM to disk when creating an lpp_source took 304 MB of disk space during our tests. If you really need to save disk space and therefore do not want to copy the contents of the AIX V4.1 Server CD-ROM to disk, you can use one of the three following methods to migrate your clients:

1. Migrate the additional LPPs using NIM cust operations:

a. Create a spot, an lpp_source pointing to the mounted AIX V4.1 Server CD and bosinst_data resources. You do not need installp_bundle resources with this method.

b. Allocate these resources to the clients, and start bos_inst operations to migrate BOS on the clients.

c. After the BOS migration of the clients is finished, create another lpp_source by copying with the `bffcreate` command the BFF images (including their requisites) from the product CD-ROM or tape to a new directory or file system on disk. This second lpp_source should contain the additional LPPs that the clients need for the migration and that are not part of the AIX V4.1 Server CD-ROM.

d. Use the second lpp_source resource and installp_bundle resources containing lists of LPPs to start NIM cust operations that install and migrate the extra LPPs on the clients. NIM cust operations perform software customizations on running NIM clients; they install LPPs or PTFs on top of BOS.

Advantage: The set up is easy and does not require scripts.

Disadvantage: Two steps are required: bos_inst and cust NIM operations.

2. Migrate the additional LPPs with NIM customization scripts:

a. Create a new directory or file system, and fill it with the BFF images of the additional LPPs (including requisites) that need to be migrated, by using the `bffcreate` command from the product tape or CD-ROM. Do not define this directory as an lpp_source, but just export it with NFS.

b. Create script resources in NIM that mount this directory over NFS, and use the `installp` command to install the extra LPPs from it. These scripts are executed on the clients after the BOS migration if you allocate them to the clients together with the other resources before starting the bos_inst operation. You will probably need several scripts for different clients because LPPs are not the same for all clients.

The scripts should contain commands like:

```
mount <NIM_master_hostname>:<directory_with_extra_lpps><client_mount_point>
installp -agX -d <client_mount_point> fileset1 fileset2 fileset3 ...
```

If needed, requisite filesets are installed and file systems are extended. Make sure that the mount points exist on the clients before mounting.

c. Create a spot, an lpp_source pointing to the mounted AIX V4.1 Server CD-ROM and bosinst_data resources. You do not need installp_bundle resources with this method.

d. As the last step, allocate the resources to the clients, and start the bos_inst operations using the lpp_source on the mounted AIX V4.1 Server CD-ROM. First, BOS is migrated; then, during the NIM customization phase, the scripts are executed on the clients, and the extra LPPs are migrated, too.

Advantage: It is a one-step NIM operation (bos_inst).

Disadvantage: The set up is complex and requires long NIM customization times on clients.

3. Migrate additional LPPs via /etc/firstboot on the clients:

a. Create a spot, bosinst_data resources and two lpp_source resources. One lpp_source points to the AIX V4.1 server CD-ROM. The other

lpp_source is created by copying the BFF images of the extra LPPs, including the requisites, to a new directory or file system on disk by using the bffcreate command. No installp_bundle is needed when using this method.

b. Create script resources for the clients that echo the appropriate nimclient commands to /etc/firstboot so that the additional LPPs get installed when the clients boot for the first time. /etc/firstboot should contain commands like:

```
nimclient -o allocate -a lpp_source=<lpp_source_with_extra_LPPs>
nimclient -o cust -a filesets="fileset1 fileset2 fileset3 ..."
```

Because of the LPP lists, you will probably need specific scripts for each client. The scripts basically allocate the second lpp_source containing the extra LPPs when the clients first reboot after the BOS migration. Then they start a pull installation of the extra LPPs initiated from the clients.

c. Allocate all resources, including spot, bosinst_data, two lpp_source resources, and the script resources to the clients, and start the NIM bos_inst operations. The BOS migrations will start on the clients. After the first reboot, the LPP migrations will take place.

Advantage: It is a one-step NIM operation (bos_inst).

Disadvantage: The set up is complex and requires long boot times during the first boot of the clients.

Attention!

These three migration methods require a complex setup, and the migration cannot be performed in one step. Therefore, you should use one of these methods only if absolutely necessary. Also, we did not test these methods in detail. The intent here is to show you, in principle, what you could do if you wanted to save disk space.

6.8.8 Installing Fixes During Migration

Few LPPs do not need to be migrated and just require some PTFs to be able to work after the migration. To find out details about this type of LPPs, please refer to *A Holistic Approach to AIX V4.1 Migration, Planning Guide, SG24-4651.*

It is possible to install fixes during the migration process with NIM. The following method can be used to accomplish this (the NIM master must be at AIX V4.1.3 level or later):

1. Copy the BFF fix images into the lpp_source directory before defining the lpp_source resource.
2. Create fix_bundle files containing lists of fixes that need to be installed, for example:

```
IX53674
IX54156
```

In AIX V4.1, fixes are usually refered to by Authorized Program Analysis Report (APAR) numbers.

3. Define fix_bundle resources together with all other resources that point to the locations of the fix_bundle files.

4. Allocate the fix_bundle resources together with all other resources to the clients where you want to install the fixes during the migrations.
5. Start the migrations with the bos_inst operations.

Using this technique, NIM installs the fixes on the clients that have allocated fix_bundle resources during the migration. You can check from the NIM master if the fixes have been installed with the NIM fix_query operation.

We will not go into more detail in our scenario because only very few LPPs need this kind of migration. But with the steps given above, you should be able to install fixes if needed.

6.8.9 Additional Migration Tasks

Some LPPs may need additional migration tasks after their installation in AIX V4.1 to work with their AIX V3.2 configuration. If these tasks can be handled via shell scripts, you will be able to run them on the clients using NIM script resources.

To run shell scripts on clients with NIM, you have to define script resources that point to the path names of the scripts stored in the file systems on the servers. You can run these scripts at three different times:

1. During BOS installation

   If you allocate the script resources to clients before you start the bos_inst operations, the scripts will be executed during the BOS installation on the clients after all the `installp` operations have been performed but before the first reboot occurs.

2. During cust operations with LPPs

   You can run scripts during cust operations which install extra LPPs on clients. In this case, the script resources must be allocated to the clients in addition to the lpp_source you require and the installp_bundle which is optional. The scripts will be executed during the cust operations after all `installp` commands have been performed.

3. With extra cust operations

   You can also just run scripts on clients without installing anything. To achieve this, you only allocate script resources to the clients, and run the cust operations without additional parameters, for example:

   ```
   # nim -o cust <NIM client name>
   ```

We will not go into more detail in our scenario because it is very difficult to give general instructions for this kind of task. But with the details given above, you should be able to run scripts on clients, if needed.

Sometimes you have to take care of complex migration tasks that you do not want to handle with scripts or with tasks that are not identical on many clients. In these cases, manual intervention may be necessary.

6.8.10 Installing Non-AIX Software

Non-AIX software can only be distributed with NIM if it uses the `installp` command for installation (for example, CATIA). If this is the case, you can place the software into the lpp_source resource with the `bffcreate` command, and add it to the installp_bundle resources where necessary. However, it depends on the software whether it migrates its configuration when installed over an earlier version of itself or if it simply overwrites everything.

If the software is installed with other commands, like `tar` for example, or via special procedures, you should check with your software vendor to see if the vendor has a procedure for the migration.

6.8.11 Testing One Client First

First, you should always test your chosen migration method and the NIM setup with one client. When the migration is finished and you see after testing the client's functions that everything works as expected, you can migrate the other clients over the LAN. If the migration on the first client did not work as expected, you still have time to fix problems before you migrate all of your clients. This prevents spreading problems into your whole environment.

6.8.12 Migrating Gateways

If you have several networks in your environment which are linked via gateways, you have to be careful about the migration of the gateways.

It is not possible to migrate clients reached via a gateway and to migrate the gateway at the same time. During the migration of a gateway, the network reached via this gateway is not accessible. Therefore, you should follow this procedure:

1. Migrate the clients in the network where the NIM master is located, including the gateways connected to this network.
2. Migrate the clients in the networks that are reached via the gateways migrated in the first step, including other possible gateways connecting networks in which you have not yet migrated any clients.
3. If you have more networks that are reached via gateways migrated in the previous step, migrate the clients in these networks now, including other possible gateways connecting networks in which you have not yet migrated any clients.
4. Repeat step three if you have more networks.

You can compare this procedure with the structure of an onion that consists of shells included inside each other. First, you migrate the innermost shell, then the next shell, then the next one, until you reach the outermost shell.

When the gateways are migrated, the TCP/IP definitions are kept. Therefore, all the gateways should still work after the migration without additional operations.

6.8.13 Planning Duration

In this section, we give you some durations that we experienced during our testing. You can use them to plan the duration of your migration preparation and the migration itself.

Durations we experienced:

- Configuration of a NIM master that is not a server including installation:

  1/2 day

- Configuration and set up of a NIM master that is also a complete server:

  1 day

- Configuration and set up of additional servers:

  1/2 day

- Migration of clients:

  2-3 hours, depending on the hardware and the amount of software

Set up means, in this case, that all the resources are defined and in the case of spot and lpp_source, populated with files.

The minimum time you have to plan for just setting up one NIM master that also serves all resources is one day. Therefore, we think that the migration of clients via NIM only pays if you have a minimum of three clients. If you have fewer clients, it will be quicker to migrate them with local media.

6.9 Setting Up NIM for Migration: Flowchart

In this section, a flowchart shows you all the required steps for setting up NIM for an unattended migration of multiple AIX V3.2 systems to AIX V4.1.

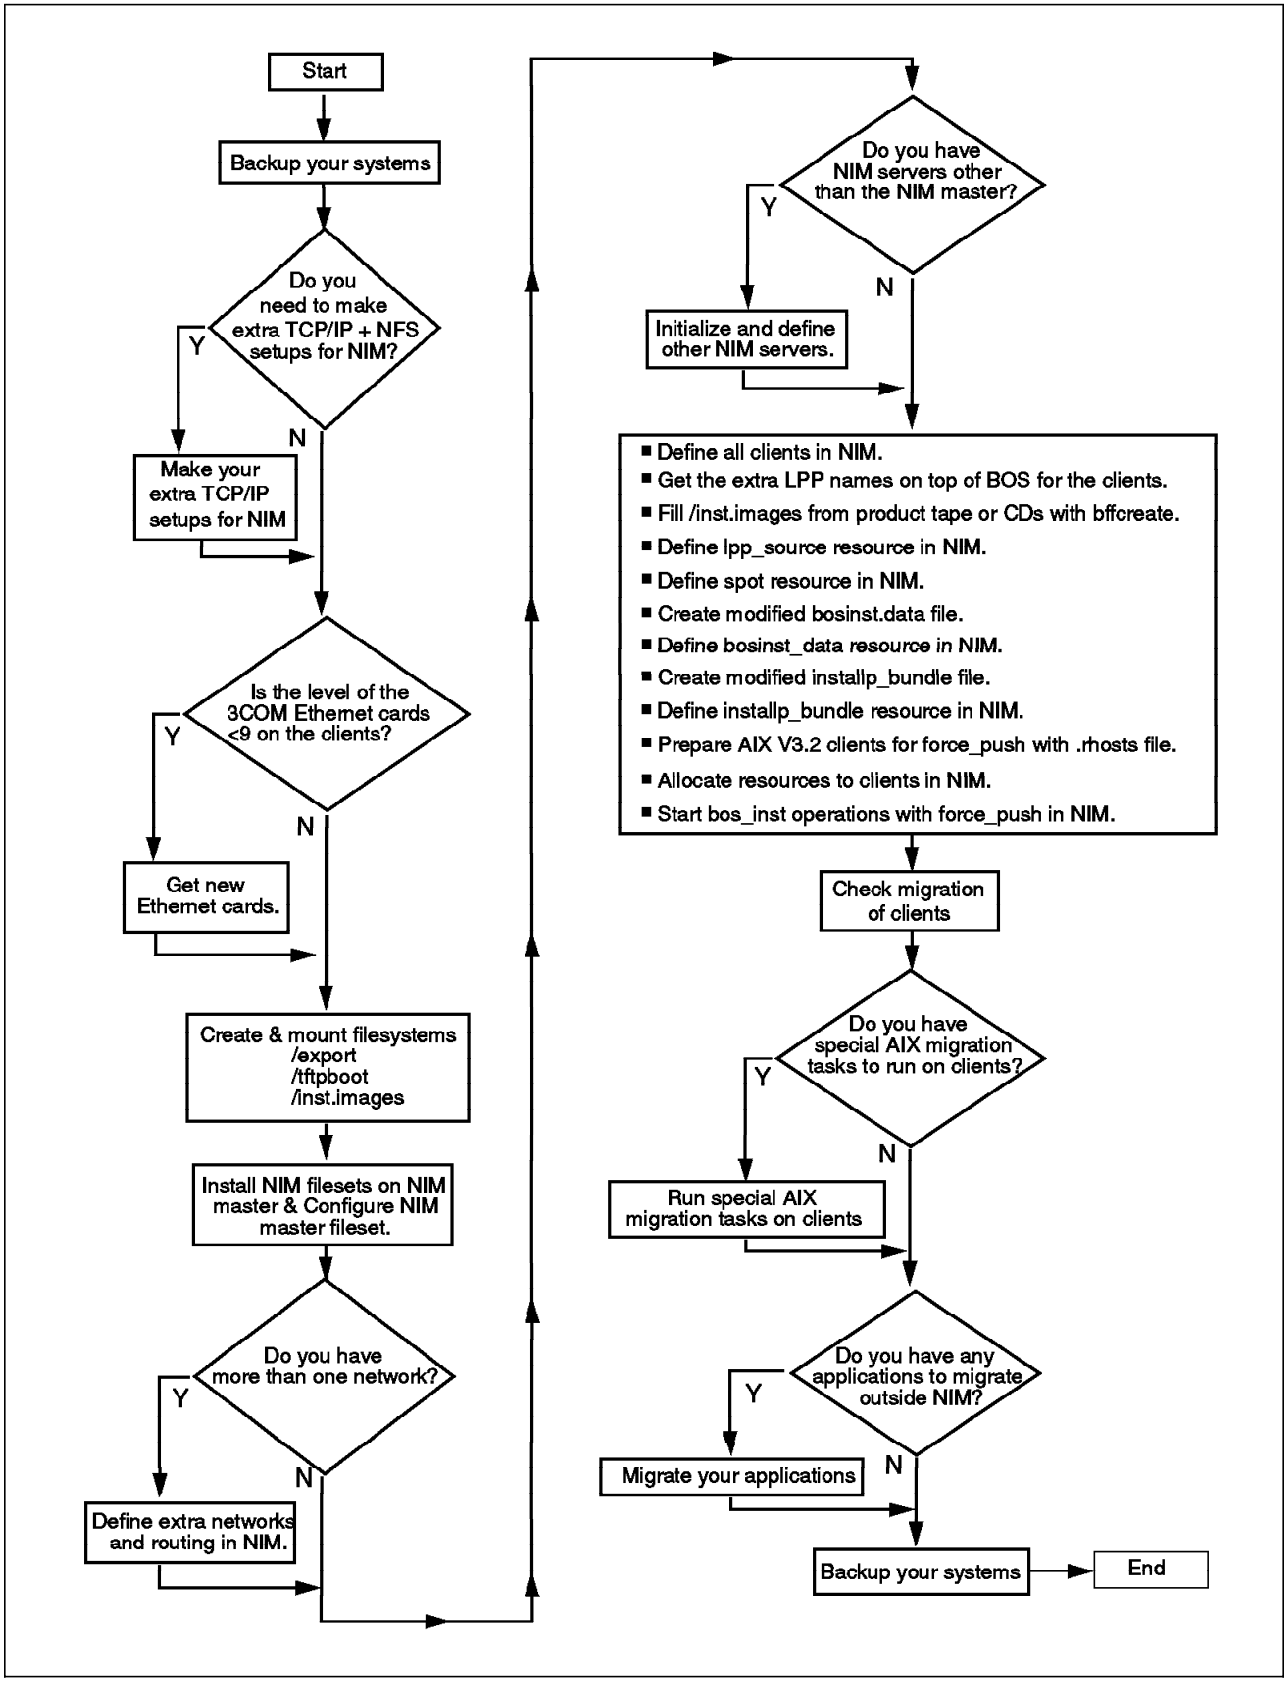

Figure 11. NIM Migration Setup Steps

We used the recommended methods we developed in the previous sections. If you want to use other methods, you need to work with a different flow that is adapted to your needs.

6.10 Setting Up NIM for New Installation

Setting up NIM for a new installation is very similar to setting up NIM for a migration.

In fact, there are only two differences:

1. The bosinst_data resource

   The bosinst_data resource must point to a different bosinst.data file than the one used for migration. Depending on which installation method is used for the new installation, the bosinst.data file must contain either the line

   ```
   INSTALL_METHOD = overwrite
   ```

   or the line

   ```
   INSTALL_METHOD = preserve
   ```

   For details concerning installation methods, please refer to the section 6.2, “Choosing the Installation Method” on page 131.

   The stanza describing the target disks should also be edited to reflect the disk drive configurations of the target systems:

   ```
   target_disk_data:
       LOCATION =
       SIZE_MB =
       HDISKNAME = hdisk0
   ```

   If you leave it blank, as you can do during the migration, the disks chosen for the new installation will not always be predictable.

   Note: This stanza has a decreasing priority, meaning that LOCATION overrides SIZE_MB, which overrides HDISKNAME.

   All other lines can be left unmodified, as compared to the migration setup which is described in section 6.12, “Migration Scenario” on page 159.

2. The source for the bos_inst operation

   As we described in section 6.8, “Setting Up NIM for Migration: General Considerations” on page 145, it is generally possible to choose either the spot or the lpp_source or the mksysb resource as the source for BOS during the bos_inst operation. For migrations, however, only the lpp_source resource can be chosen.

   For new installations, you can either use the spot or the lpp_source resource. For the differences, please refer to section 6.8, “Setting Up NIM for Migration: General Considerations” on page 145. Using the mksysb resource is described in section 6.11, “Setting Up NIM for Cloning” on page 158 because it does more than just a new installation.

6.11 Setting Up NIM for Cloning

Setting up NIM for cloning means to set it up to install the clients from a mksysb image. You have to take into account the following points:

- The mksysb resource

  A mksysb resource has to be set up and allocated to the clients. It is simply a pointer to a file that contains the mksysb image you want to install. NIM only supports AIX V4.1 mksysb images. This mksysb resource is chosen as the source for the bos_inst operation.

  Attention!

  Even if NIM is not using the spot and lpp_source resources for installation, they must still be created and allocated to the clients during cloning. The spot resource is needed because the boot images required for the network boots are generated during the spot creation. The lpp_source is needed as a source for additional software, like the NIM client filesets, TCP/IP filesets and the devices filesets, even if they are in the mksysb image.

- The bosinst.data files

  The bosinst.data files must be set up if you want to run unattended installations at the clients. You can use two sources for the bosinst.data files.

  1. The mksysb image

     The mksysb image that you are using for installation can contain a modified bosinst.data file in the / directory. It is the same for all the clients.

  2. The bosinst_data resource

     If you allocate bosinst_data resources to the clients, the bosinst.data file from the mksysb image will be overridden. This way, you can allocate special bosinst.data files to some clients which cannot use the bosinst.data file in the mksysb image for some reason.

  The following lines must be set up in the bosinst.data files for unattended cloning:

```
control_flow:
    CONSOLE = /dev/lft0   -> or /dev/tty0 for ASCII terminals
    PROMPT = no
    EXISTING_SYSTEM_OVERWRITE = yes
    TCB = no              -> Note that TCB must be enabled in the mksysb
                             image for a value of yes to be meaningful
                             here.
target_disk_data:
    LOCATION =
    SIZE_MB =
    HDISKNAME = hdisk0
```

- The installp_bundle resource

  The installp_bundle resource is not needed during cloning with NIM. The amount of software that gets installed on the clients is determined by the mksysb image.

6.12 Migration Scenario

In this section, we describe the scenario used in our tests. We show the prerequisites that must be met and detailed steps on how to migrate this scenario environment with NIM.

6.12.1 Scenario Description

Our scenario consists of one NIM master and three NIM clients. Two clients are connected through a thin (coaxial cable) Ethernet LAN to the master. One client is connected to the token-ring LAN that is connected to the Ethernet over a gateway. In most of our tests, the NIM master is at the same time the server for all the resources, and the nim_client_250 is simply one NIM client. But for some special tests, we also installed the nim_client_250 from the product media and integrated it into the NIM environment later to use it as a server in addition to the NIM master.

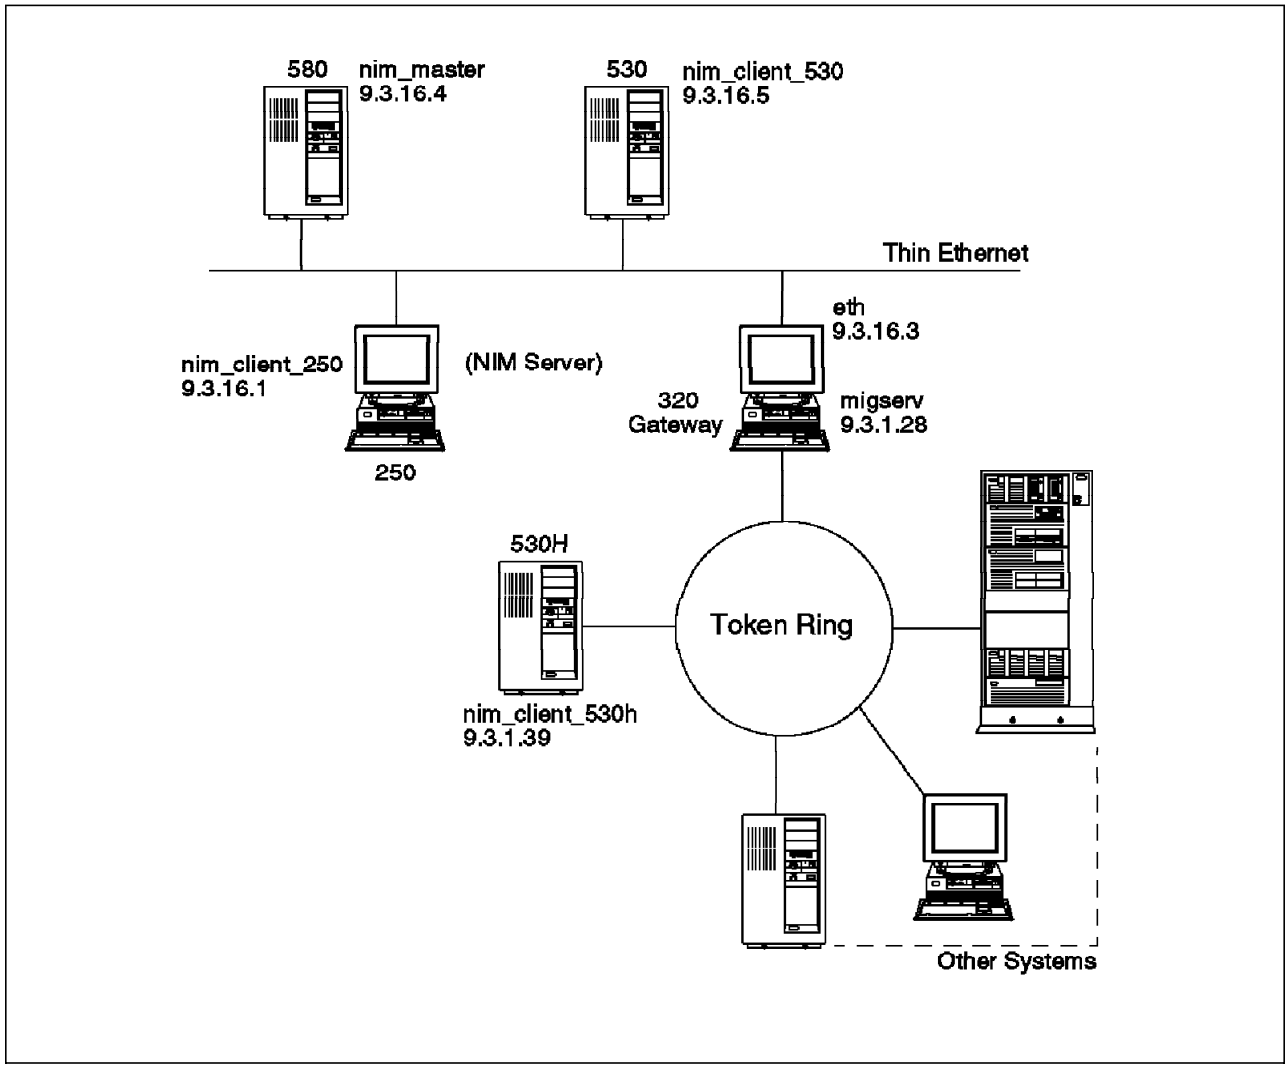

Figure 12. Our NIM Scenario

Even if the scenario is not as big as a real environment, it is sufficient to develop methods and steps which are necessary to migrate multiple systems on a LAN without local interventions.

The flowchart in section 6.9, “Setting Up NIM for Migration: Flowchart” on page 155 illustrates the steps required to migrate the entire environment.

Attention!

You must be the root user for all the steps that follow.

6.12.2 Backing Up the Systems

Remember, the first step should always be to back up your systems. You need it to be able to recover your environment as quickly as possible in case of problems.

We either use `mksysb` tapes to back up our clients, or we put the `mksysb` images on our AIX V4.1 NIM master over NFS and re-install the clients with the AIX V3.2 Network Installation Server mechanism. Please refer to section 6.3, “Backing Up AIX V3.2 Systems” on page 134, for details. We do not describe the backup procedure in detail because we assume that you already have a procedure in place for making your backups.

6.12.3 Checking Network Communication

The next steps are used to check the network communication.

6.12.3.1 TCP/IP and NFS

You have to make sure that the complete TCP/IP setup for your environment is done before you set up NIM. This includes definitions of networks, initialization of gateways, name serving, and routing. All the host names for the systems participating in the NIM environment (master, servers and all clients) must be resolvable.

Here is an excerpt of the file /etc/hosts we used on the NIM master to resolve the host names:

```
9.3.16.4        nim_master                          # NIM master
9.3.16.1        nim_client_250                      # NIM client Ethernet
9.3.16.5        nim_client_530                      # NIM client Ethernet
9.3.1.39        nim_client_530h                     # NIM client token ring
9.3.16.3        eth.itsc.austin.ibm.com     eth     # gateway ethernet side
9.3.1.28        migserv.itsc.austin.ibm.com migserv # gateway token ring side
```

NFS must also run on the system planned to be the NIM master. If you are unsure about this, you can check with the command `lssrc -g nfs` for active subsystems, and use the fast path `smitty mknfs` to start NFS. We use the ASCII version of SMIT, called `smitty` , because it is faster.

We do not go into more detail about TCP/IP and NFS here because we assume that your network and servers are already set up since you were already using them before the migration.

6.12.3.2 Ethernet Card Levels on the Clients

If you use an Ethernet LAN, you should check the level of the 3COM Ethernet cards with the command:

```
# lscfg -v
```

This does not apply to integrated Ethernet modules on system planars.

The field called ROS Level and ID must be 9 or higher.

> **Attention!**
>
> If you have earlier levels, you will see a blinking 888 during the network boot on the client and get the codes 103-203-280 when pressing the reset button.

Get newer Ethernet cards if this is the case.

6.12.4 Preparing the NIM Master and NIM Servers

This section describes the steps to prepare the NIM master and NIM servers.

6.12.4.1 Creating File Systems

In this step, we create file systems for some of the NIM resources for easier administration:

- /export

  Use the /export file system for directories that contain the spot, the bosinst_data and the installp_bundle resources. A size of 150 MB should be sufficient for one spot resource.

  The name /export is arbitrary; this means that you can use another file system name. However, by convention, the spot resource is located under /export/exec. If you use the name /export and you plan to serve the spot resource with the NIM master, it is best to create the file system before the installation of the NIM filesets. During the installation of the bos.sysmgt.nim.master fileset, some subdirectories are created in the /export directory that you would have to recreate if you create the file system after the master is configured.

- /tftpboot

  During the creation of the spot resource, the network boot images are generated and put into the /tftpboot directory. If you do not use an extra file system for this directory, it is created in the / file system. NIM creates several network boot images for different system platforms (rs6k, rs6ksmp or rspc) and for different network interfaces (tok, ent and fddi). Together, they need about 20 MB of disk space, which would enlarge the / file system considerably.

  Boot images are always created in the /tftpboot directory. If you create an extra file system for these images, the file system must be mounted on a directory called /tftpboot.

- /inst.images

  This file system is created for the lpp_source resource. As all software that the clients need for migration is stored here, it can grow very big. The minimum is more than 300 MB, but it can grow to more than 1 GB. Therefore, you should even consider creating an extra volume group for it.

  The name /inst.images is arbitrary; this means that you can use another file system name. However, by convention, /inst.images is used to store BFF images for LPPs.

To create the three file systems, you can specify only one block as the size. This way, they are only created with the size of one Physical Partition (PP). NIM

takes care of the necessary increase during the creation of the resources. This is a way to avoid creating file systems that are bigger than necessary.

You can use the SMIT fast path smitty crjfs to create the file systems. You can also use the path:

```
# smitty
System Storage Management (Physical & Logical Storage)
-> File System
   -> Add / Change / Show / Delete File Systems
      -> Journaled File Systems
         -> Add a Journaled File System
```

Following is a sample of a SMIT screen showing the creation of the /export file system in the rootvg:

```
                          Add a Journaled File System

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

                                                      [Entry Fields]
  Volume group name                                    rootvg
* SIZE of file system (in 512-byte blocks)            [1]                        #
* MOUNT POINT                                         [/export]
  Mount AUTOMATICALLY at system restart?               yes                       +
  PERMISSIONS                                          read/write                +
  Mount OPTIONS                                        []                        +
  Start Disk Accounting?                               no                        +
  Fragment Size (bytes)                                4096                      +
  Number of bytes per inode                            4096                      +
  Compression algorithm                                no                        +

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

You should make sure that the file systems are mounted automatically after system restart and that they are in read/write mode. You can also use the crjfs command directly from the command line with the correct options to create the file systems.

Please do not forget to mount the file systems after the creation. The output of your mount command should look similar to this:

```
  node       mounted        mounted over    vfs       date        options
-------- ---------------  ---------------  ------ ------------ ---------------
         /dev/hd4         /                jfs    Dec 05 16:46 rw,log=/dev/hd8
         /dev/hd2         /usr             jfs    Dec 05 16:46 rw,log=/dev/hd8
         /dev/hd9var      /var             jfs    Dec 05 16:46 rw,log=/dev/hd8
         /dev/hd3         /tmp             jfs    Dec 05 16:46 rw,log=/dev/hd8
         /dev/hd1         /home            jfs    Dec 05 16:47 rw,log=/dev/hd8
         /dev/lv00        /export          jfs    Dec 05 16:47 rw,log=/dev/hd8
         /dev/lv01        /tftpboot        jfs    Dec 05 16:47 rw,log=/dev/hd8
         /dev/lv02        /inst.images     jfs    Dec 12 15:38 rw,log=/dev/loglv00
```

If you plan to use servers other than the NIM master, the /export and /tftpboot file systems should be created on the servers serving the spot resources. The file system, /inst.images, should be created on the servers serving the lpp_source resources.

6.12.4.2 Installing NIM Filesets

In this step, we install the NIM filesets on the master. As we plan to use the NIM master also as the server, the following filesets must be installed:

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

You can use the path:

```
# smitty
Software Installation and Maintenance
-> Install and Update Software
   -> Install/Update Selectable Software (Custom Install)
      -> Install Software Products at Latest Level
         -> Install New Software Products at Latest Level
```

or the fast path `smitty install_latest` to get the following SMIT screen:

```
                  Install Software Products at Latest Level

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

                                                      [Entry Fields]
* INPUT device / directory for software               /dev/rmt0.1
* SOFTWARE to install                                [4.1.4.0  Network Insta> +
  PREVIEW only? (install operation will NOT occur)    no                      +
  COMMIT software updates?                            yes                     +
  SAVE replaced files?                                no                      +
  ALTERNATE save directory                           []
  AUTOMATICALLY install requisite software?           yes                     +
  EXTEND file systems if space needed?                yes                     +
  OVERWRITE same or newer versions?                   no                      +
  VERIFY install and check file sizes?                no                      +
  Include corresponding LANGUAGE filesets?            yes                     +
  DETAILED output?                                    yes                     +



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

The fileset bos.sysmgt.nim.client is installed automatically as a prerequisite.

Note: If you are working with CD-ROM, be sure to use an AIX V4.1 Server CD-ROM because the NIM master fileset bos.sysmgt.nim.master is not part of the Client CD-ROMs.

6.12.4.3 Configuring NIM Master Fileset

This step shows the configuration of the NIM master fileset. The system where the NIM master fileset is installed and configured becomes the unique NIM master system that controls the whole NIM environment. From this system, all the following set up and installation operations are initiated.

With the fast path `smitty nim`, you can jump directly to the NIM main menu in SMIT. From there, you can reach all NIM menus in SMIT.

```
                        Network Installation Management

Move cursor to desired item and press Enter.

  Configure Network Installation Management Master Fileset
  Manage Network Objects
  Manage Machine Objects
  Manage Resource Objects
  Create IPL ROM Emulation Media

F1=Help             F2=Refresh          F3=Cancel           F8=Image
F9=Shell            F10=Exit            Enter=Do
```

You can use the fast path `smitty nimconfig` , which is the same as the path:

```
# smitty
Software Installation and Maintenance
-> Network Installation Management
   -> Configure Network Installation Management Master Fileset
```

to get to the following SMIT screen:

```
            Configure Network Installation Management Master Fileset

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

                                                      [Entry Fields]
* Network Object Name                                 [net1]
* Primary Network Install Interface                   [en0]                    +
* Port Number for Network Install Communications      [1058]                    #
    Ring Speed                                        []
    Cable Type                                         bnc                     +









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

If you prefer to work from the command line, you can also use the following command:

```
# nimconfig -a pif_name=en0 \
            -a master_port=1058 \
            -a netname=net1 \
            -a cable_type=bnc
```

Here, the network object, net1, that represents the Ethernet to which most systems are connected, is created. You have to specify a primary network install interface. This is the network interface the NIM master will use to communicate with the network. It must be defined because it could lead to confusion if the master had several interfaces and it was not clear which one should be used for the communication with the NIM environment.

The TCP/IP port number, 1058, is the default for the NIM network communication between the NIM master and its clients. You can change the port number in case of conflicts. Check the /etc/services file for potential conflicts before selecting a port number.

The cable type must be chosen for Ethernet LANs. For token-ring LANs, the ring speed must be specified instead.

6.12.4.4 Extra Network Setup

In this step, we define the second network object, net2, that represents the token-ring. One of our clients is connected to the token-ring network. The token-ring is connected to the Ethernet via a gateway. The NIM master which is located in the Ethernet must know how to reach the client in the token-ring. Therefore, you must first define a second network and then the routing between the networks.

To define the second network, you must define a network object in NIM. You can use the fast path `smitty nim_mknet`, which is the same as the path:

```
# smitty
Software Installation and Maintenance
-> Network Installation Management
   -> Manage Network Objects
      -> Define a Network Object
```

and select the object type, **tok**, to get to the following SMIT screen:

```
                              Define a Network Object

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

                                                        [Entry Fields]
 * Network Object Name                                  [net2]
 * Network Object Type                                   tok
 * Network IP Address                                   [9.3.1.0]
 * Subnetmask                                           [255.255.255.0]
   Comments                                            []

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

If you prefer to work from the command line, you can also use the following command:

```
# nim -o define \
      -t tok \
      -a net_addr=9.3.1.0 \
      -a snm=255.255.255.0 \
      net2
```

Note: You can specify the address of any host on the subnet being defined and NIM will calculate the IP address of the network.

We define a second network, named net2, of the type token-ring. The network IP address consists of the first three digits of the client address in the token-ring because we use a subnetmask of 255.255.255.0.

Now you have to define the NIM routing. You can use the fast path smitty nim_mkroute, which is the same as the path:

```
# smitty
Software Installation and Maintenance
-> Network Installation Management
   -> Manage Network Objects
      -> Manage Network Install Routing
         -> Define a Network Install Route
```

and choose the **originating** and the destination network to get the following SMIT screen:

```
                        Define a Network Install Route

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

                                                      [Entry Fields]
* Originating Network Object Name                     net1
* Gateway Used by Originating Network                [eth]
* Destination Network Object Name                     net2
* Gateway Used by Destination Network                [migserv]
  Force                                               no                      +









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

If you prefer to work from the command line, you can also use the command:

```
# nim -o change \
      -a routing1='net2 eth migserv' \
      net1
```

Specify the names for the originating and for the destination network and the IP host names that are used for both sides of the gateway.

6.12.4.5 Setting Up Other NIM Servers

If you want to achieve the best possible performance during the migration, you can use NIM servers other than the NIM master to serve some resources.

Of the four resource types we use (lpp_source, spot, bosinst_data and installp_bundle), the lpp_source type is the most critical concerning performance. You have to choose the lpp_source as the source for the software for the migration. Therefore, it is accessed during the whole migration over NFS from the clients to serve the software products they need for installation. Serving lpp_source resources from several servers at a time can improve your migration times significantly for big networks with many clients. For details on sizing, please refer to section 6.6.4, “Performance and Sizing” on page 142.

In the case of migrations, the spot resource type is mainly used during the remote network boot of the clients. If you want to boot many clients at a time, it can also make sense to set up several servers for spot resources. However, the influence on the overall performance of the migration is lower compared to the lpp_source type.

The two remaining resource types, bosinst_data and installp_bundle, are only used at certain points of the migration to determine parameters or for some setup work. They are not accessed all the time, and therefore their influence on performance is small. This also applies to script resources if you use them.

A good indicator for the influence on performance is the amount of data that is transmitted over the network to the clients. You can deduct this amount from the size of the resources on disk. The lpp_source is the biggest, followed by the spot resource. The other resources are very small compared to them.

If you want to set up NIM servers other than the NIM master, these servers must first be stand-alone NIM clients running AIX V4.1. You have two possibilities to achieve this:

1. Install or migrate to AIX V4.1 systems planned to be NIM servers, using NIM.

   If you plan to set up the NIM master also as a complete server with all resources, you can use the NIM master first to install or migrate your planned server systems. The advantage in this case is that systems installed or migrated via NIM are automatically set up to be NIM clients. Therefore, you can reach your server systems with NIM afterwards without additional setup work. With this method, the flow to upgrade your environment looks like this:

   a. Configure the NIM master, which is also a server.

   b. Set up all resources on the master.

   c. Use the master to install or migrate other planned server systems.

   d. Set up the resources on the other NIM servers.

   e. Migrate your clients using the NIM master and servers.

   Attention!

   You need to make sure that in addition to the NIM master, the NIM servers are also able to resolve the client's host names. This may not be configured automatically during the server installation or migration with NIM.

2. Install or migrate planned NIM servers using a media.

   If you do not want to set up the NIM master as a complete server first, you will need to install or migrate your planned server systems with tape or CD-ROM. The disadvantage in this case is that you must manually set them up as stand-alone NIM clients. This requires several steps:

   a. Set up TCP/IP and NFS on the servers.

      You have to make sure that TCP/IP is configured on the servers. This includes definitions of networks, gateways, name serving, and routing.

      Attention!

      All the host names for the systems participating in the NIM environment must be resolvable on the servers, including the NIM master's host name.

      NFS must also run on the servers. If you are unsure about this, you can check with the command `lssrc -g nfs` for active subsystems, and use the fast path `smitty mknfs` to start NFS.

   b. Install the NIM client fileset bos.sysmgt.nim.client on the servers.

      You need to install the fileset bos.sysmgt.nim.client on the servers so that they can communicate with the NIM master. You can use the fast path `smitty install_latest` to do this.

c. Configure the servers as NIM clients.

   This cannot simply be done from the NIM master because the servers are not reachable yet via NIM. Instead, the configuration is initiated from the clients. They contact the NIM master and ask to be configured.

   To achieve this, you can use the fast path `smitty niminit` on the client. This is the same as the path:

```
# smitty
Software Installation and Maintenance
-> Network Installation Management
   -> Configure Network Installation Management Client Fileset
```

   You will get to the following SMIT screen:

```
            Configure Network Installation Management Client Fileset

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

                                                      [Entry Fields]
* Machine Object Name                                [nim_client_250]
* Primary Network Install Interface                  [en0]                    +
* Host Name of Network Install Master                [nim_master]
* Port Number for Network Install Communications     [1058]                    #
    Ring Speed                                       []                        #
    Cable Type                                        bnc                     +
  Hardware Platform Type                              rs6k                    +
  IPL ROM Emulation Device                           []                        /
  Comments                                           []

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

   You can also use the following command:

```
# niminit -a name=nim_client_250 \
          -a pif_name=en0 \
          -a master=nim_master \
          -a master_port=1058 \
          -a cable_type=bnc \
          -a platform=rs6k
```

   With this SMIT screen or the `niminit` command, you can set up systems as NIM clients and tell them to contact the NIM master and to ask to be configured. Now, the NIM master can reach the systems like any other client. Therefore, resources can now be defined to be located on the systems. They can be used as NIM servers.

This second method requires less work on the NIM master but more work on the servers. You may want to use it because you may have some systems which are already running AIX V4.1 that you now want to integrate into your NIM environment for use as additional servers. If you want to use this method, the flow to upgrade your environment will look like this:

a. Install or migrate the planned server systems from AIX V4.1 media (if they are not installed already).

b. Configure the NIM master.

c. Set up the planned servers as stand-alone NIM clients.

d. Set up the resources on the master.

e. Set up the resources on the servers.

f. Migrate your clients using the NIM master and servers.

Even if this section does not seem to fit into the flow, we have to give you the information at this time because it influences the next steps, for example where to put the resources. In all the following sections on resources, we describe what has to be done to put the resources on servers other than the NIM master.

Attention!

If you put lpp_source resources on servers, please remember that you have to copy the BFF images on them first with bffcreate.

6.12.4.6 Defining NIM Clients

This step defines the clients nim_client_250, nim_client_530 and nim_client_530h that we want to migrate. You can use the fast path smitty nim_mkmac or the path:

```
# smitty
Software Installation and Maintenance
-> Network Installation Management
   -> Manage Machine Objects
      -> Define a Machine Object
```

and choose the **hardware** and **machine type** to get the SMIT screen:

```
                          Define a Machine Object

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

                                                    [Entry Fields]
* Machine Object Name                              [nim_client_250]
  Hardware Platform Type                            rs6k
  Machine Object Type                               standalone
  Primary Network Install Interface
    Network Object Name                             net1
*   Host Name                                       [nim_client_250]
    Network Adapter Hardware Address                [0]
    Network Adaptor Logical Device Name             []
*   Cable Type                                       bnc                    #
  IPL ROM Emulation Device                          []                     /
  CPU Id                                            []
  Comments                                          []




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

You can also use the following command:

```
# nim -o define \
      -t standalone \
      -a platform=rs6k \
      -a if1='net1 nim_client_250 0' \
      -a cable_type1=bnc \
      nim_client_250
```

With this SMIT screen or command, define the nim_client_250 which is connected to the Ethernet (network object net1). You must specify the platform type (rs6k, rs6ksmp or rspc) and the machine type (can be diskless, dataless or stand-alone). The rs6k means standard IBM RISC System/6000; the rs6ksmp means IBM RISC System/6000 SMP models, and rspc means PCI-based systems with the PowerPC processor.

For migration, you only use rs6k and stand-alone. You must also tell NIM the IP host name of the system. For easier administration, it is better to use the same name for the NIM object name and the IP host name. This is not necessary; the NIM object name is arbitrary and is only used by NIM.

You also need to specify the cable type if using Ethernet or the token-ring speed if using token-ring.

Note: The hardware address of the network adapter always needed to be specified before AIX V4.1.3; starting with AIX V4.1.3, a “0” can be specified instead.

The other details are either not needed, or are obtained by NIM itself later.

To define the client, nim_client_530h, in the token-ring (NIM network object net2), the SMIT screen looks like this:

```
                              Define a Machine Object

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

                                                      [Entry Fields]
* Machine Object Name                                [nim_client_530h]
  Hardware Platform Type                              rs6k
  Machine Object Type                                 standalone
  Primary Network Install Interface
    Network Object Name                               net2
*   Host Name                                         [nim_client_530h]
    Network Adapter Hardware Address                  [0]
    Network Adaptor Logical Device Name               []
*   Ring Speed                                        [16]                    #
  IPL ROM Emulation Device                            []                      /
  CPU Id                                              []
  Comments                                            []




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

The NIM command you can use is:

```
nim -o define \
    -t standalone \
    -a platform=rs6k \
    -a if1='net2 nim_client_530h 0' \
    -a ring_speed1=16 \
    nim_client_530h
```

The only differences compared to the definition of the nim_client_250 are the network object, net2, and the ring speed parameter instead of the cable type.

6.12.4.7 Getting Additional LPP Names

This step is a preparation for the creation of the lpp_source and installp_bundle resources that will be used to migrate the clients.

To be able to migrate the clients in one step, all the software must be put into the lpp_source, and installp_bundle resources listing the additional software that must be installed on top of BOS, must be created. Therefore, you have to get a list of the extra LPPs that must be migrated in addition to BOS.

Usually, that is simply the software you ordered additionally to your BOS license. It should therefore be fairly easy to find out the additional LPPs. In our case, the list consists of the following LPPs: SNA, HCON and XLC.

6.12.4.8 Populating /inst.images

This step describes how to put all the software that is necessary for the lpp_source resource into the directory /inst.images.

The lpp_source will be used as the source for the client migration. That means that all the software that the clients need for the migration, including all the requisites, must be put in here. Also, NIM needs a set of required images (the support images) to define the directory /inst.images as the lpp_source. For a list of the support images, please refer to 6.8.7.1, “Working with Product Tape” on page 149.

To reduce the probability for a requisite problem or a NIM problem, you can choose the easiest way here. Just copy everything that is on the product media to the /inst.images directory. This sacrifices some disk space on the server, but avoids a possible source for errors. If you do not want to do it this way, you can use the list of LPPs from the previous step in addition to the list of required images to populate the lpp_source.

Attention!

Be sure to use your server product media for this step. Not all necessary software products may be included in the client media.

Working with product tape, you can use the fast path `smitty bffcreate` which is the same as:

```
# smitty
Software Installation and Maintenance
-> Install and Update Software
   -> Copy Software to Hard Disk for Future Installation
```

to get this SMIT screen:

```
                Copy Software to Hard Disk for Future Installation

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

                                                      [Entry Fields]
* INPUT device / directory for software               /dev/rmt0.1
* SOFTWARE package to copy                           [all]                    +
* DIRECTORY for storing software package             [/inst.images]
  DIRECTORY for temporary storage during copying     [/tmp]
  EXTEND file systems if space needed?                yes                     +

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

You can also use the following command:

```
# bffcreate -d /dev/rmt0.1 -t /inst.images -X all
```

When using SMIT or the bffcreate command to load all the software products from the product tape to the directory /inst.images, it is important to specify that the file system /inst.images can be extended if space is needed because we only created it with a size of 4 MB. During the bffcreate, it will be extended to the necessary size. The bffcreate process can take some hours, depending on the amount of software that needs to be loaded and the speed of the media.

If you work with CD-ROM, you must first run the bffcreate command with the AIX V4.1 Server CD-ROM, choosing /dev/cd0 as the input device. Then you must repeat it with all your product CD-ROMs. Usually, you get a separate CD-ROM for every product you ordered. If you got several independent product tapes, the same applies. Multi-volume product tapes will be processed by bffcreate in one pass.

If you work with servers other than the NIM master for the lpp_source resources, you will start the bffcreate on the servers using the servers' tape or CD-ROM drives to load the software products to the /inst.images file systems you created on the servers.

When you are finished, all the software you need for migration should be in the directory /inst.images. You can list the BFF images with the ls command. There is also a file called .toc created by the bffcreate command. This is the table of contents describing the loaded software products. It shows the names of the BFF images, the individual filesets with their requisites and more. You can examine this file with your favorite ASCII editor if you are unsure if all needed filesets are there.

6.12.4.9 Defining the lpp_source Resource

This step defines the lpp_source resource. The lpp_source is pointing to the directory /inst.images that contains all the BFF images we already loaded. It will be used by NIM as the source for the software packages sent over the network to the clients during migration.

You can use the fast path `smitty nim_mkres`, which is the same as the path:

```
# smitty
Software Installation and Maintenance
-> Network Installation Management
   -> Manage Resource Objects
      -> Define a Resource Object
```

and choose the **lpp_source resource type** to get to the following SMIT screen:

```
                          Define a Resource Object

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

                                                      [Entry Fields]
* Resource Object Name                               [images]
* Resource Type                                       lpp_source
* Server of Resource                                 [master]                +
* Location of Resource                               [/inst.images]           /
  Source of Install Images                           []                      +/
  Names of Option Packages                           []
  Comments                                           []

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

You can also use the following command:

```
# nim -o define \
      -t lpp_source \
      -a location=/inst.images \
      -a server=master \
      images
```

The object name, images, is arbitrary and can be chosen to be describe the content of the resource. The server and the location of the resource specify the system and the directory where the resource is located.

If a system other than the NIM master is used as the server for an lpp_source resource, you will use its NIM object name to specify the server parameter. The location you will specify in that case will be the path name of the directory on the server where you loaded the software products with `bffcreate`. You can also copy the BFF images over from another server or from the master with the `rcp`

command if you set up more than one lpp_source resource. The command to define the lpp_source resource must always be started on the NIM master.

Attention!

We do not specify the source parameter for the creation of the lpp_source because we want to use a directory that already contains the necessary BFF images as the location.

This way, NIM only checks if the necessary support images are there and defines the lpp_source pointing to the directory. You should not get any errors during the definition of the lpp_source if you loaded all the software products from your server product media to the directory /inst.images in the previous step.

If nothing is missing, the simages attribute will be set for the lpp_source. If you get errors, NIM will usually tell you which support images are missing. NIM cannot check for software products other than the support images because it does not know what the clients have installed that needs to be migrated.

You can load missing support images into the directory /inst.images by using bffcreate the way we did it in the previous step. When this is done, NIM must be told to check the lpp_source again for the missing support images.

This is accomplished by the NIM check operation. You can use the fast path smitty nim_res_op or the path:

```
# smitty
Software Installation and Maintenance
-> Network Installation Management
   -> Manage Resource Objects
      -> Perform Operations on Resource Objects
```

and choose the **images resource** and the **check operation** to get this SMIT screen:

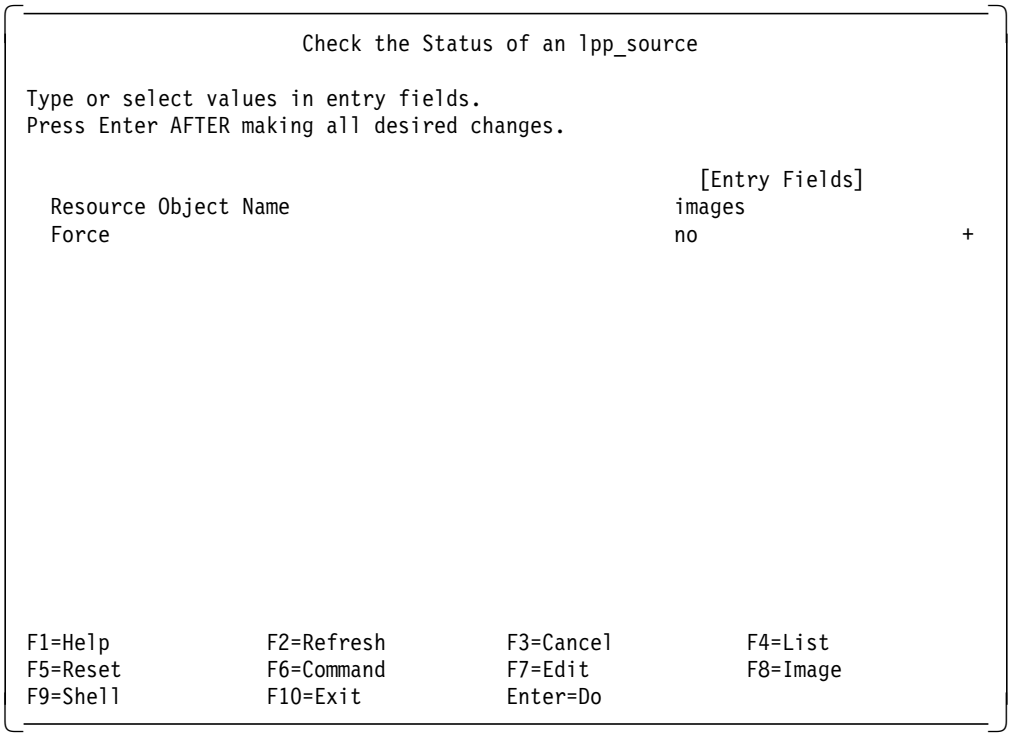

You can also use the following command:

```
# nim -o check images
```

The check operation will give you a list of missing support images if some images are still missing, or it will set the simages attribute for the lpp_source if everything is there.

To be sure that the simages attribute is set for your lpp_source, you can run the following command after the definition of the lpp_source or after the check operation:

```
# lsnim -l images
```

It gives you the following output:

```
images:
   class       = resources
   type        = lpp_source
   alloc_count = 0
   server      = master
   location    = /inst.images
   Rstate      = ready for use
   prev_state  = unavailable for use
   simages     = yes
```

The line

```
   simages     = yes
```

is present if the simages attribute is set.

You can repeat the definition of lpp_source resources, including the preparation with bffcreate, several times to create different lpp_source resources on different servers for performance reasons. Of course, the resource names must be different. Alternatively, you can copy an lpp_source by following the

lpp_source definition process and specifying an existing lpp_source as the ″source″ for the new lpp_source.

6.12.4.10 Defining the Spot Resource

In this step, we create the spot resource. During the spot creation, the boot images for the network boot of the clients are generated.

First, you need to create the directory where the spot resource will be located:

```
# mkdir /export/exec
```

Then you can use the fast path `smitty nim_mkres`, which is the same as the path:

```
# smitty
Software Installation and Maintenance
-> Network Installation Management
   -> Manage Resource Objects
      -> Define a Resource Object
```

and choose the **spot resource type** to get to the following SMIT screen.

```
                          Define a Resource Object

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

                                                      [Entry Fields]
 * Resource Object Name                              [spot1]
 * Resource Type                                      spot
 * Server of Resource                                [master]                +
 * Source of Install Images                          [images]                +
 * Location of Resource                              [/export/exec]           /
   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           F10=Exit            Enter=Do
```

The following command achieves the same:

```
# nim -o define \
      -t spot \
      -a location=/export/exec \
      -a server=master \
      -a source=images \
      -a auto_expand=yes \
      spot1
```

Again, the resource name is arbitrary and can be chosen to best describe the content of the resource. The server and the location of the resource specify the system and the directory where the resource is located.

If a system other than the NIM master is used as the server for a spot resource, you will use its NIM object name to specify the server parameter. In that case, the location you specify will be the path name of the directory on the server where the spot resource will be created. For the creation of spot resources on other servers, the source for the install images must be local media or lpp_source resources with the simages attribute set. No remote media is allowed. If lpp_source resources are used, they can be anywhere in the NIM environment. The command to define the spot resource must always be started on the NIM master.

In our case, where the NIM master is also the server, we specify the lpp_source resource called images as the source of the install images. This is the fastest way to create the spot because only disk-to-disk operations are involved. You can also specify the CD-ROM or tape devices with appropriate product media loaded as sources, but this takes longer. Therefore, it makes sense to first create the lpp_source and then the spot.

We specify the directory /export/exec for the location of the spot resource. It must exist when the spot is created, or you will get an error message. The actual path name that is used for the spot is /export/exec/spot1/usr.

We could use the /usr file system of the master or servers to create the spot. But because of the reasons given in section 6.8.4, “BOS Installation Operation Sources” on page 146, we prefer to use an extra file system, /export.

Again, it is important that we allow that the file systems can be expanded if space is needed because we only created the /export and /tftpboot file systems with a size of one PP.

If you want to check the state of the spot resource you created, you can use the command `lsnim -l <spot name>`. For the spot named spot1, we get the following output:

```
spot1:
   class        = resources
   type         = spot
   alloc_count  = 0
   server       = master
   location     = /export/exec/spot1/usr
   Rstate       = ready for use
   prev_state   = unavailable for use
   version      = 04
   release      = 01
   if_supported = rs6k ent
   if_supported = rs6k fddi
   if_supported = rs6k tok
   if_supported = rs6ksmp ent
   if_supported = rs6ksmp tok
   if_supported = rspc ent
   if_supported = rspc tok
```

You should get a similar output for your spot resources.

You can repeat this step several times to create different spot resources on different servers for performance reasons. Of course, the resource names must be different. Again, existing spots can be used as the source to create other spots in the NIM ennironment, effectively copying the spot to a new resource.

6.12.4.11 Creating the bosinst.data File

In this step, we create a modified bosinst.data file. The NIM bosinst_data resource that is built from this file will be used to specify the installation parameters for the clients. Without this resource, an unattended installation of the clients is not possible.

Probably, you will need several different bosinst.data files and resources for different client types. For example, clients with different console types (graphical and ASCII) need different bosinst.data files. Or clients with or without the need to initialize the TCB during installation also need different files. Also, if you plan to install clients with a New or Complete Overwrite Install instead of migrating them, you will need different bosinst.data files.

Therefore, we created the directory /export/bosinst_datas as the repository for all the bosinst.data files:

```
# mkdir /export/bosinst_datas
```

To this directory, we first copy the bosinst_data template which is located in /var/adm/ras:

```
# cp /var/adm/ras/bosinst.data /export/bosinst_datas/bosinst.data.mig.lft.tcb
```

You can name the file bosinst.data.mig.lft to make clear that it is a bosinst.data file that will be used for the migration of clients with Low-Function Terminal (LFT) (graphical) consoles and set the Trusted Computing Base (TCB) initialization to yes.

Now, use an ASCII editor, such as vi, to edit the bosinst.data file. In the first part of the file, the different parameters and possible values are explained. You can use the following setup for your migrations:

```
control_flow:
    CONSOLE = /dev/lft0
    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 = yes
    INSTALL_TYPE = full
    BUNDLES =

target_disk_data:
    LOCATION =
    SIZE_MB =
    HDISKNAME =

locale:
    BOSINST_LANG = en_US
    CULTURAL_CONVENTION = en_US
    MESSAGES = en_US
    KEYBOARD = en_US
```

The following values are important:

- CONSOLE = /dev/lft0

Specifies the device that will be the client's console. For ASCII terminals, CONSOLE = /dev/tty0, or another TTY should be used.

- INSTALL_METHOD = migrate

  Specifies that the installation method will be a migration.

- PROMPT = no

  Indicates that no prompting for values at the clients should be used, if possible. It works only if all other required values are present. If you see three zeros and a rotating line at the migration start, all required values are found.

- EXISTING_SYSTEM_OVERWRITE = yes

  Specifies that an existing system can be overwritten.

- TCB = yes

  Indicates that the TCB should be initialized during migration. This was not possible before AIX V4.1.3. If the TCB should not be initialized, TCB = no should be used.

- locale:

  ```
  BOSINST_LANG = en_US
  CULTURAL_CONVENTION = en_US
  MESSAGES = en_US
  KEYBOARD = en_US
  ```

  Specifies the locales that will be used during installation and the messages that will be installed.

Please note that for migrations, the target disks do not need to be defined because the migration will automatically choose the disks where the rootvg is located.

Attention!

If your bosinst.data files miss important parameters or do not match with client configurations, the installation process reverts to prompting at the client consoles. You should watch the client consoles carefully during the first migrations to make sure that you do not miss any prompts.

6.12.4.12 Defining the bosinst_data Resource

In this step, we use the customized bosinst.data file we created in the last step to define the bosinst_data resource in NIM.

You can use the fast path `smitty nim_mkres`, which is the same as the path:

```
# smitty
Software Installation and Maintenance
-> Network Installation Management
   -> Manage Resource Objects
      -> Define a Resource Object
```

and choose the **bosinst_data resource type** to get to the following SMIT screen:

```
                          Define a Resource Object

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

                                                      [Entry Fields]
* Resource Object Name                               [mig_param_lft_tcb]
* Resource Type                                       bosinst_data
* Server of Resource                                 [master]                 +
* Location of Resource                               [/export/bosinst_datas/>  /
  Comments                                           []

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

You can also use the following command:

```
# nim -o define \
      -t bosinst_data \
      -a server=master \
      -a location=/export/bosinst_datas/bosinst.data.mig.lft.tcb \
      mig_param_lft_tcb
```

Again, the resource name is arbitrary and can be chosen to best describe the content of the resource. The server and the location of the resource specify the system and the path name where the bosinst.data file is located.

If a system other than the NIM master is used as the server for a bosinst_data resource, you will use its NIM object name to specify the server parameter. The location you will specify in that case will be the path name of the bosinst.data file on the server that is already created there. The command to define the bosinst_data resource must always be started on the NIM master.

You can repeat this step several times to create different bosinst_data resources. Of course, their resource names must be different.

6.12.4.13 Creating the installp_bundle File

In this step, we create the installp_bundle file. This file will be used for the definition of the installp_bundle resource in NIM. It consists of a list of software products that need to be migrated on the clients in addition to BOS.

During the Migration Install, BOS, with its accessory products, like TCP/IP or X-Windows, is migrated automatically. Extra software products you ordered, like HCON or SNA for example, are not migrated automatically. You have to list them in the installp_bundle file to be able to run the whole migration of the clients in one step. Since clients do not have exactly the same software configuration, you will probably need different lists.

In an installp_bundle file, you can specify the LPP names, the package names or the fileset names:

1. LPP names

   If the LPP names are specified, all packages with all filesets belonging to the packages are installed.

2. Package names

   If the package names are specified, all filesets belonging to the packages are installed.

3. Fileset names

   If the fileset names are specified, these specific filesets are installed.

All the software gets installed with the necessary requisites. Therefore, it can only be installed if the requisites are present in the lpp_source resource.

If too many software products are installed on the clients, the products can always be de-installed later. The NIM maint operation can be used for that purpose.

The best way to get the exact LPP, package and fileset names you need, is to look into the TOC from the lpp_source that you already created on disk or to access the TOC of the product media.

To look at the TOC from the lpp_source, you can just use an ASCII editor:

```
# vi /inst.images/.toc
```

If you must look at the TOC of your product media, it works differently for product CD-ROMs and tapes:

- Product CD-ROMs

  Product CD-ROMs have a directory structure on them. The BFF images and the TOC are stored in the directory /usr/sys/inst.images.

  In order to read the TOC, you must first define a CD-ROM file system using, for example, the fast path `smitty crcdrfs`.

  Then you must mount the CD-ROM onto its mount point. Now you can look at the TOC by using an editor, for example:

  ```
  # vi /<CD_mount_point>/usr/sys/inst.images/.toc
  ```

- Product tapes

  On product tapes, the third image on the first tape is the TOC. You need first to copy it to the disk before you can look at it.

  Be sure that your tape device works with a block size of 512 bytes. You can check this with the following command:

  ```
  # lsattr -l rmtx -E
  ```

  where rmtx is your tape unit. Check for the line:

  ```
  block_size    512  BLOCK size (0=variable length)   True
  ```

  If the block size is different, you can change it, for example, with the fast path `smitty chgtpe`.

  To read the TOC to disk, you can use the command:

  ```
  # dd if=/dev/rmtx of=/tmp/TOC fskip=2
  ```

where /dev/rmtx is your tape unit and /tmp/TOC is the TOC file to be created.

Now you can look at the TOC using an editor, for example with:

```
# vi /tmp/TOC
```

Here is a small excerpt of a TOC to show you what you need to look for:

```
01:295:28442624 4 R I xlC.C++ {
xlC.C++.browser 03.01.0001.0000 01 N B en_US C Set ++ for AIX Browser
[
*prereq xlC.C++.cmp 3.1.1.0
*prereq X11.base.rte 4.1.0.0
*prereq bos.net.tcp.client 4.1.0.0
*prereq xlC.rte 3.1.1.0
%
/usr/lpp/xlC/browser 8432
/usr/lib/objrepos 8
/usr/lib/X11/app-defaults 8
/usr/bin 8
INSTWORK 1232 768
%
xlCbrs.obj 2.99.0.0
%
%
%
]
xlC.C++.cmp 03.01.0001.0000 01 N U en_US C Set ++ for AIX Compiler
[
*prereq xlC.C 3.1.1.0
%
```

It describes the BFF image, xlC.C++ which is the image number, 295, on the first volume. We can see this in the line:

```
01:295:28442624 4 R I xlC.C++ {
```

The package name is the same as the BFF image. Therefore, we now know that the package name is xlC.C++. The LPP name is always the first portion of the package name. Therefore xlC is the LPP name. Sometimes package and LPP name are equal (hcon for example).

If you want to use single filesets, then you should look for lines like:

```
xlC.C++.browser 03.01.0001.0000 01 N B en_US C Set ++ for AIX Browser
xlC.C++.cmp 03.01.0001.0000 01 N U en_US C Set ++ for AIX Compiler
```

Here, the two filesets, xlC.C++.browser and xlC.C++.cmp, are listed with their Version Release Modification Fix (VRMF) numbers and text descriptions.

You have to prepare the installp_bundle file for the LPPs you got in section 6.12.4.7, “Getting Additional LPP Names” on page 172. In our case, we use the same LPPs for all clients: SNA, HCON and XLC.

We use the LPP names instead of the package or fileset names for our installp_bundle file because this is the safest way to list all packages belonging to an LPP without forgetting any packages. Using the TOC, as described above, we got the three LPP names that we need to put into our installp_bundle file: hcon, sna and xlC. Now, you can create the file.

First create an extra directory for the installp_bundle files:

```
# mkdir /export/bundles
```

Then, use an ASCII editor to create an installp_bundle file called leftovers1:

```
# vi /export/bundles/leftovers1
```

The contents of the file is the following:

```
hcon
sna
xlC
```

6.12.4.14 Defining the installp_bundle Resource

In this step, we define in NIM the installp_bundle resource that is used to start the migration of the LPPs in addition to BOS. We use the installp_bundle file we created in the last step.

You can use the fast path `smitty nim_mkres`, which is the same as the path:

```
# smitty
Software Installation and Maintenance
-> Network Installation Management
   -> Manage Resource Objects
      -> Define a Resource Object
```

and choose the **installp_bundle resource type** to get to the following SMIT screen:

```
                          Define a Resource Object

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

                                                      [Entry Fields]
* Resource Object Name                               [leftovers1]
* Resource Type                                       installp_bundle
* Server of Resource                                 [master]                 +
* Location of Resource                               [/export/bundles/leftov>  /
  Comments                                           []

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

You could also use the following command:

```
# nim -o define \
      -t installp_bundle \
      -a server=master \
      -a location=/export/bundles/leftovers1 \
      leftovers1
```

Again, the resource name is arbitrary and can be chosen to best describe the content of the resource. The server and the location of the resource specify the system and the path name where the installp_bundle file is located.

If a system other than the NIM master is used as the server for an installp_bundle resource, you will use its NIM object name to specify the server parameter. The location you will specify in that case will be the path name of the installp_bundle file on the server that is already created there. The command to define the installp_bundle resource must always be started on the NIM master.

You can repeat this step several times to create different installp_bundle resources. Of course, their resource names must be different.

Now, all the resources that are needed by NIM for the migrations are created.

6.12.5 Preparing the Client

In order to be able to push-boot the AIX V3.2 clients, one preparatory step must be done on them. Each client must have a .rhosts file in the / directory to allow the NIM master root access to the clients.

The contents of the /.rhosts files should therefore look like this:

```
<nim_master_hostname> root
```

In our example, the /.rhosts file looks like this:

```
nim_master root
```

The files can be copied to the client by using the File Transfer Protocol (FTP) or another mechanism.

Attention!

Please remember that all clients must be running and be reachable through the network to initiate an unattended migration with NIM.

The key switches of all clients must be set to Normal before the bos_inst operation is performed because of the force_push attribute we used.

6.12.6 Starting the Migration

Now we can start the migration of the clients from the NIM master. Two actions have to be done for this.

6.12.6.1 Allocating Resources

In this step, we allocate the resources needed for migration to the clients. The resources are: lpp_source, spot, bosinst_data, and installp_bundle.

You can use the fast path `smitty nim_alloc`, which is the same as the path:

```
# smitty
Software Installation and Maintenance
-> Network Installation Management
   -> Manage Machine Objects
      ->  Manage Network Install Resource Allocation
          -> Allocate Network Install Resources
```

and choose the **nim_client_250 machine object** to get to the following SMIT screen:

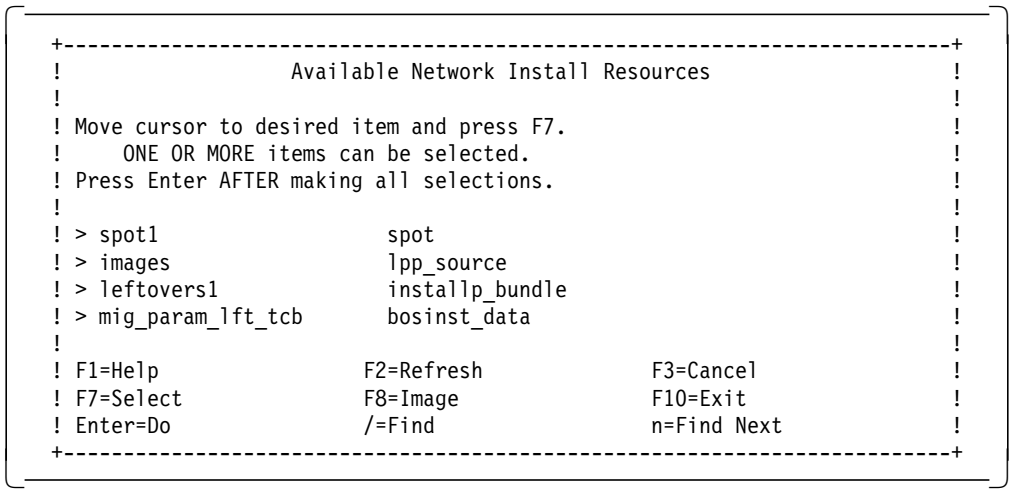

Use **F7** to select all the resources for the allocation.

You can also work with the command:

```
#
nim -o allocate \
      -a spot=spot1 \
      -a lpp_source=images \
      -a installp_bundle=leftovers1 \
      -a bosinst_data=mig_param_lft_tcb \
      nim_client_250
```

You can always check later if the resources are allocated with commands like the following:

```
# lsnim -c resources nim_client_250
```

The output should look like this:

```
spot1                  spot
images                 lpp_source
leftovers1             installp_bundle
mig_param_lft_tcb      bosinst_data
```

During the bos_inst operation which we will start in the next step, you may see other resources, like boot resources for example, that are allocated to the clients. NIM is doing this automatically.

You have to repeat this step for all the clients. NIM only allows the allocation of resources to one client at a time with one allocate operation. If you have many clients, we recommend that you write a small shell script to allocate the resources. It should contain something like the following:

```
# nim -o allocate \
    -a spot=spot1 \
    -a lpp_source=images \
    -a installp_bundle=leftovers1 \
    -a bosinst_data=mig_param_lft_tcb \
    nim_client_250
#
echo "\nAllocated Resources for NIM client nim_client_250:\n"
```

```
lsnim -c resources nim_client_250
#
echo "\nSleep 5 Seconds\n"
sleep 5
#
nim -o allocate \
    -a spot=spot1 \
    -a lpp_source=images \
    -a installp_bundle=leftovers1 \
    -a bosinst_data=mig_param_lft_tcb \
    nim_client_530
#
echo "\nAllocated Resources for NIM client nim_client_530:\n"
lsnim -c resources nim_client_530
#
echo "\nSleep 5 Seconds\n"
sleep 5
#
nim -o allocate \
    -a spot=spot1 \
    -a lpp_source=images \
    -a installp_bundle=leftovers1 \
    -a bosinst_data=mig_param_lft_tcb \
    nim_client_530h
#
echo "\nAllocated Resources for NIM client nim_client_530h:\n"
lsnim -c resources nim_client_530h
```

The delay of five seconds with the sleep 5 command between the different allocate operations is precautionary because we must prevent NIM from allocating resources to clients too fast in succession. This could lead to errors while exporting the resources with NFS.

All resources are deallocated when the bos_inst operation is finished, even if it failed. That means that you have to reallocate all four resources to the clients when you restart the bos_inst operation because of problems.

All allocate operations must always be started from the NIM master, even if the resources are located on other servers in the network.

6.12.6.2 Starting bos_inst Operations

In this step, we start the bos_inst operations on the clients. The migration of BOS and additional LPPs on the clients will take place in one step because of our preparations. All bos_inst operations are always started on the NIM master.

To start the migration on the nim_client_250, you can use the fast path smitty nim_mac_op, which is the same as the path:

```
# smitty
Software Installation and Maintenance
-> Network Installation Management
   -> Manage Machine Objects
      ->  Perform Operations on Machine Objects
```

and choose the **nim_client_250 machine object** and the **bos_inst operation** to get to the following SMIT screen:

```
                          Perform a Network Install

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

                                                      [Entry Fields]
  Machine Object Name                                 nim_client_250
  Source for BOS Runtime Files                        rte                    +
  installp Flags                                      [-agX]
  Remain NIM client after install?                    yes                    +
  Initiate Boot Operation on Client?                  yes                    +
  Set Boot List if Boot not Initiated on Client?      no                     +
  Force Unattended Installation Enablement?           yes                    +

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

You can also use the following command:

```
# nim -o bos_inst \
      -a source=rte \
      -a installp_flags=-agX \
      -a force_push=yes \
      nim_client_250
```

Specifying rte as the source for the BOS runtime files means that the lpp_source is used as the software source for the migration. This is the only possible choice.

The specified flags for the installp command indicate that the software is applied with requisites and that file systems on the client can be extended if necessary.

Setting the unattended installation enablement to yes means switching the force_push attribute to yes. This is the only way to push-boot AIX V3.2 clients. The force_push attribute tells NIM that the target of the bos_inst operation is a machine that does not necessarily have the bos.sysmgt.nim.client fileset installed and configured. Accordingly, NIM will attempt to NFS mount or copy the minimal client support to the target system to perform an unattended installation or migration of the base operating system. Use of the force_push attribute requires that the key on the client is in the Normal position.

The other parameters can be left in their default states.

You have to repeat this step for all the clients. NIM only allows the BOS installation or migration of one client at a time with a bos_inst operation. If you have many clients, we recommend that you write a small shell script to start the migrations. It should contain something like the following:

```
echo "\n Starting bos_inst operation for nim_client_250...\n"
#
nim -o bos_inst \
    -a source=rte \
    -a installp_flags=-agX \
    -a force_push=yes \
    nim_client_250
#
echo "\n Starting bos_inst operation for nim_client_530...\n"
#
nim -o bos_inst \
    -a source=rte \
    -a installp_flags=-agX \
    -a force_push=yes \
    nim_client_530
#
echo "\n Starting bos_inst operation for nim_client_530h...\n"
#
nim -o bos_inst \
    -a source=rte \
    -a installp_flags=-agX \
    -a force_push=yes \
    nim_client_530h
```

You should not start several of these scripts simultaneously because NIM is allocating more resources while starting the bos_inst operation. You must prevent NIM from allocating resources to clients simultaneously. This could lead to errors while exporting the resources with NFS. Using the top-down flow of one script, NIM has enough time to avoid difficulties.

In this step, NIM reboots the clients and starts the migration. When the BOS and LPP installations are finished, the clients are rebooted again and connected to the network as NIM clients. After the completion of the bos_inst operations, all resources will be automatically deallocated.

6.12.6.3 If Something Goes Wrong

If something goes wrong during the migration, you may have to reset the state of the clients with a NIM reset operation to be able to deallocate the resources and work on them if necessary. Resources can only be changed if they are not allocated.

To reset the clients, you can use the fast path `smitty nim_mac_op`, which is the same as:

```
# smitty
Software Installation and Maintenance
-> Network Installation Management
   -> Manage Machine Objects
      -> Perform Operations on Machine Objects
```

and choose the **client** (for example the nim_client_250) and the **reset operation** to get to the following SMIT screen:

```
                        Reset the NIM State of a Machine

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

                                                      [Entry Fields]
  Machine Object Name                                nim_client_250
  Force                                              yes                     +

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

You can also use the following command:

```
# nim -o reset -a force=yes nim_client_250
```

We usually had to set the force parameter to yes to be able to reset the clients.

After this step, you are able to deallocate the resources. You can use the fast path `smitty nim_dealloc`, which is the same as:

```
# smitty
Software Installation and Maintenance
-> Network Installation Management
   -> Manage Machine Objects
      -> Manage Network Install Resource Allocation
         -> Deallocate Network Install Resources
```

and choose the **machine object** (for example the nim_client_250) to get to the following SMIT screen:

```
  +--------------------------------------------------------------------------+
  !                   Allocated Network Install Resources                    !
  !                                                                          !
  ! Move cursor to desired item and press F7.                                !
  !     ONE OR MORE items can be selected.                                   !
  ! Press Enter AFTER making all selections.                                 !
  !                                                                          !
  ! > spot1                   spot                                          !
  ! > images                  lpp_source                                    !
  ! > leftovers1              installp_bundle                               !
  ! > mig_param_lft_tcb       bosinst_data                                  !
  !                                                                          !
  ! F1=Help                F2=Refresh             F3=Cancel                  !
  ! F7=Select              F8=Image               F10=Exit                   !
  ! Enter=Do               /=Find                 n=Find Next                !
  +--------------------------------------------------------------------------+
```

As you can see, we use **F7** to choose all four resources to be deallocated.

You can also use the command:

```
# nim -o deallocate \
      -a spot=spot1 \
      -a lpp_source=images \
      -a installp_bundle=leftovers1 \
      -a bosinst_data=mig_param_lft_tcb \
      nim_client_250
```

Now, it is possible to make changes to the resources. Remember to run the NIM check operation on the lpp_source resources if you copy additional software into them (see section 6.12.4.9, “Defining the lpp_source Resource” on page 174 for details). Then the resources can be allocated again to the clients, and the bos_inst operations can be restarted.

If you want to modify or delete resources and get an error message that they are locked, you can do the following:

```
# stopsrc -s nimesis
# startsrc -s nimesis
```

This stops and restarts the NIM nimesis subsystem. The resources should not be locked anymore after its restart so that you should be able to manipulate them.

6.12.7 Checking for a Successful Migration

In this section, we check if the migrations have been successful. All these steps have to be executed on the clients. For example, telnet can be used to get to the clients and run the checks.

6.12.7.1 Software Inconsistencies

First, we suggest that you use the `lppchk -v` command to check for inconsistencies in the software. It should give you no output at all if everything is fine. If it finds inconsistencies, it will give you a list of the missing software.

6.12.7.2 Operating System Level

Then check for the level of the operating system. The `oslevel` command shows you the system level. The output should look like 4.1.3.0 or 4.1.4.0, depending on the system level you use. If you get something like <4.1.4.0, this indicates that some software products are not at 4.1.4.0 level. You can produce a listing of these products with `oslevel -l <level>`.

6.12.7.3 Software Product Levels

With the command `lslpp -l` , you can check for all the software product levels. The output shows the level for each fileset so that you are able to see what has been migrated and what may still need to be migrated.

6.12.7.4 Free Disk Space

It is also a good idea to check for the free disk space after the migration is finished. This can be done by using the `df -k` command. The option `-k` is used because the disk space is reported in 512 byte blocks in AIX V4.1 by default.

6.12.8 Performing Specific Migration Tasks

If the above checks indicated that all software is migrated, you are ready with the NIM migration tasks. If anything is left over to migrate, perhaps because the software has not been in the lpp_source, you can put it into the lpp_source, and start NIM cust operations to install the missing software to the clients.

If you have any special migration tasks to perform on some clients, such as copying or modifying configuration files, you can use NIM script resources and cust operations to run these tasks. You can also run these scripts before this time by running them during the NIM bos_inst operation. Please refer to section 6.8.9, “Additional Migration Tasks” on page 153 for details.

6.12.9 Migrating Non-AIX Software

If you have to migrate non-AIX software, please refer to section 6.8.10, “Installing Non-AIX Software” on page 154, for details.

6.12.10 Backing Up the Migrated Systems

When all your migration tasks are finished, you should make backups of your systems. In AIX V4.1, you can use the fast path `smitty mksysb` to make system backups with the `mksysb` command.

You can use local media (tape drives) to make your backups. But as you work in a LAN environment, you can also store your backups on the servers via NFS. In that case, you can use NIM to reinstall the backups on the clients in case of problems. You can find details about the installation of `mksysb` images on clients with NIM in section 6.11, “Setting Up NIM for Cloning” on page 158.

After making the system backups, the migration tasks are finished.

Part 4. High Availability Cluster Multi-Processing

Chapter 7. HACMP Introduction and Terminology

Since its introduction in June of 1992, High Availability Cluster Multi-Processing (HACMP) has been a leader in high-availability technology. It is a very important product in IBM′s commercial systems strategy for UNIX environments.

This chapter will provide an introduction to the HACMP technology and define some of the relevant terminology.

7.1 Related Publications

The following publications are related directly to HACMP. For publications related to AIX in general, and AIX migration, see “Related Publications” on page xviii.

- *HACMP 4.1.1: Concepts and Facilities*, SC23-2767
- *HACMP 4.1.1: Planning Guide*, SC23-2768
- *HACMP 4.1.1: Installation Guide*, SC23-2769
- *HACMP 4.1.1: Administration Guide*, SC23-2770
- *HACMP 4.1.1: Troubleshooting Guide*, SC23-2771
- *HACMP 4.1.1: Programming Locking Applications*, SC23-2772
- *HACMP 4.1.1: Programming Client Applications*, SC23-2773
- *HACMP 4.1.1: Master Index and Glossary*, SC23-2774
- *High Availability on the RISC System/6000 Family*, SG24-4551
- *An HACMP Cookbook*, SG24-4553

7.2 High Availability Environment

The HACMP architecture enables clustered RISC System/6000 server processors to handle failures at a minimal cost. An HACMP cluster detects and recovers from failures of disks, disk adapters, networks, network adapters and processors. With the HACMP software, a cluster of loosely coupled processors, or nodes, provides application availability by transferring control from a failed system to a backup system that shares some common resources with the failed system.

The architecture provides reliable, recoverable, shared disk resources for database or online transaction processing (OLTP) transactions and client applications. This is done by a combination of HACMP software and basic AIX functions such as the Logical Volume Manager (LVM), Journalled File System(JFS) and TCP/IP. It also requires careful attention to hardware configuration. The relationship between AIX, HACMP and application software is shown in Figure 13 on page 196.

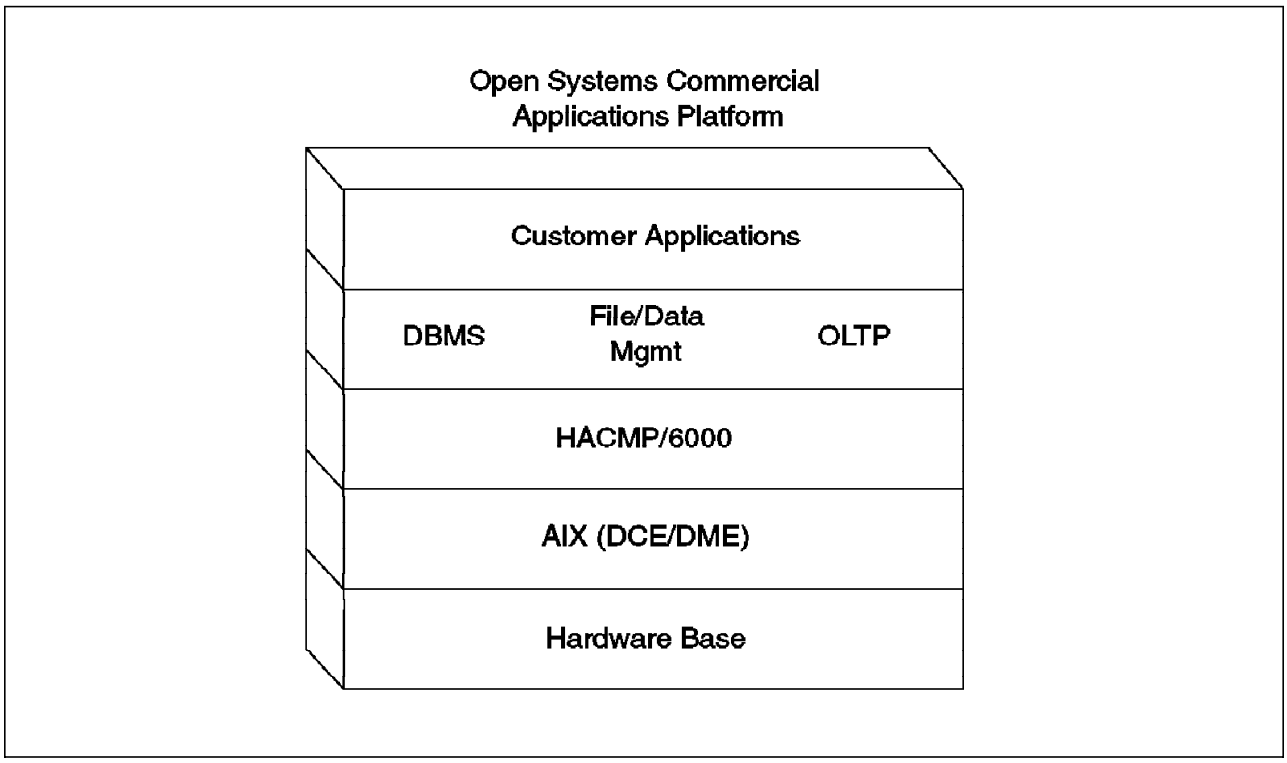

Figure 13. HACMP Model

7.3 Examples of HACMP Architecture

A sample configuration could include two processors, multiple disk drives physically connected to each of the two processors, and would utilize standard AIX/6000 Logical Volume Manager (LVM) software disk mirroring.

The typical network configuration for HACMP consists of a single or dual network for communications among the cluster nodes and clients and a point-to-point connection for communication between the cluster nodes.

An example of an HACMP cluster is shown in Figure 14 on page 197.

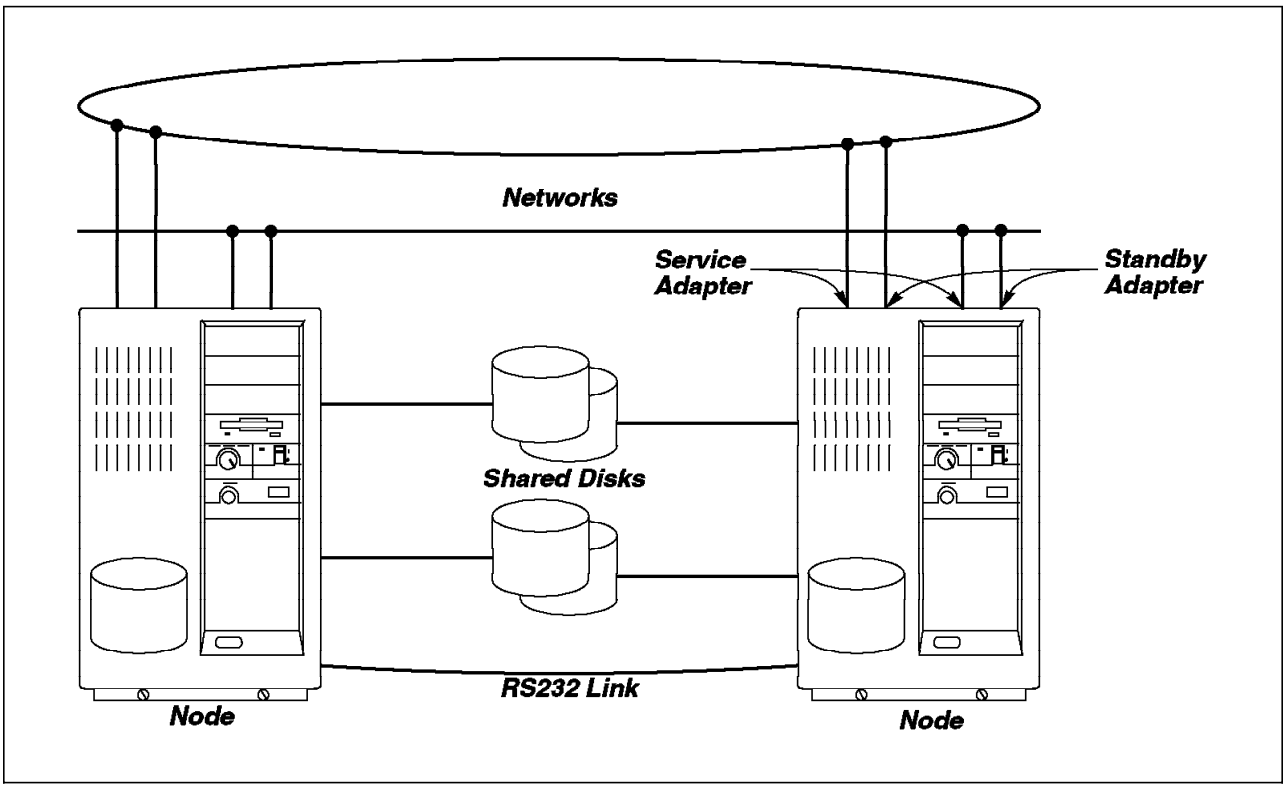

Figure 14. Example of HACMP Cluster

7.3.1 HACMP Operations

Depending on the configuration, an HACMP cluster can perform the following types of operations:

Hot Standby An active node is backed up by an idle standby node. All the shared disks and shared volume groups belong to the active node. Only after failure of this system are all those resources shifted over to the backup node.

Mutual Takeover All nodes are active, each runing its own applications using its own disk resources. In the case of a failure, the other nodes take over the failing node′s resources and applications. All applications are still available, but performance may be degraded.

Concurrent Access Two or more nodes are active simultaneously, sharing the same physical disk resources (for example, both processors could execute the same application using data on the same disk).

In earlier releases of HACMP, these three operations were referred to as Mode 1, Mode 2 and Mode 3 respectively, and were set as a global mode of operation for the entire cluster. After the release of HACMP Version 2.1, it became possible to combine these operations in a single cluster.

Note

Due to limitations in time and resources, our testing was only performed on mutual takeover operations. No testing was performed on concurrent access operations. Many of the migration considerations will apply equally to mode 3 environments, however interoperability of HACMP V3.1.1 and V4.1.1 in a concurrent access environment was *not* tested.

7.3.2 Failure Recovery

HACMP provides recovery options for the following cluster components:

Network Adapter Failure
HACMP can takeover an IP address on a standby network adapter. This feature can also include the takeover of the network adapter hardware address, ensuring that systems using the Address Resolution Protocol (ARP) will not be affected by the failure.

Network Failure
When HACMP determines that a complete network failure has occurred, all nodes in the cluster can switch IP addresses to a spare adapter on a redundant network.

Node Failure
In the case of the failure of an entire node, the resources owned by that node, can be taken over by other nodes in the cluster.

Application Failure
Application recovery is greatly dependent upon the actual applications used. HACMP provides a mechanism that allows the system administrator to monitor, stop and start applications, and to take over applications in the event of a failure.

The following recovery functions are provided directly by the AIX operating system and are complimentary to those provided by HACMP.

Disk and Disk Adapter Failure
The disk mirroring technology provided by AIX's Logical Volume Manager (LVM) allows AIX to transparently cope with the failure of a disk or disk adapter.

7.4 HACMP Terminology

This section defines some of the terminology that is used in clustered environments.

Cluster
Two or more nodes connected together using TCP/IP and each running the Cluster Manager daemon.

Cluster Manager
The main task of the Cluster Manager is to monitor the nodes and networks in the cluster for possible failures. It performs the following tasks:

- Sends and receives keepalive packets on all network interfaces
- Monitors the state of the cluster
- Executes event scripts in response to changes in cluster states

Node

An individual machine that is a member of a cluster and running the Cluster Manager.

Service Adapter

The primary connection between the node and the network. The service adapter is the interface over which clients access the services offered by the node.

Standby Adapter

A standby adapter backs up a service adapter on the same network. It can be configured to take over the IP address, and hardware address, of the service adapter.

The standby adapter is configured to be on a separate subnet to the service adapter, and in normal operations, is only used by HACMP for the exchange of keepalive or heartbeat packets.

Boot Adapter

In the situation where a node has failed and the address of its service adapter has been taken over by the standby adapter of another node, it is important that the node does not attempt to use its service address until a controlled transition back from the backup node can be made. Thus each service adapter in the cluster must be assigned a second boot address that will be used until, under the control of HACMP, the service address is acquired.

Heartbeat or Keepalive Packets

These are small packets that are exchanged regularly over every available network type between the cluster nodes. The Cluster Manager software monitors these packets to ensure that all other nodes are running correctly.

Resources

Entities that are controlled by a node and are taken over by another node in the event of a failure. Resources can include:

- Disks, volume groups, file systems
- NFS exported file systems
- IP addresses
- Applications

In the later versions of HACMP, these entities are configured into resource groups. A resource group can be either cascading, rotating, or concurrent access. Assigning a resource to one of these categories essentially establishes the failover strategy for the cluster.

- Cascading Resources

  A takeover priority for a resource is assigned to each configured cluster resource on a per-node basis. In the event of a takeover, the node with the highest priority acquires the resource. If that node is unavailable, the node with the next-highest priority acquires the resource, and so on. When a node joins a cluster and has the highest priority for a given resource, the node takes over that resource.

- Rotating Resources

  Each resource rotates among all the nodes defined in a resource chain. The node with the highest priority in the resource chain for the resource takes over for a failed node. When a failed node

rejoins a cluster, it does *not* re-acquire its resources. It comes up as a standby.

- Concurrent Access Resources

  Resources defined as concurrent access resources can be shared simultaneously by multiple nodes. All nodes concurrently accessing a resource acquire that resource when they join the cluster. There no priorities among nodes.

Event

A cluster event is a change in the cluster that the Cluster Manager detects and processes so that critical resources remain available. A cluster event can be triggered by a change affecting a network adapter, network or node, or by the cluster reconfiguration process exceeding its time limit.

System Resource Controller

HACMP deamons are controlled by the System Resource Controller (SRC). You can use the `startsrc` and `stopsrc` command to start or stop the daemons.

Journalled File System

Historically, UNIX file system recovery took a long time and file system integrity was not guaranteed. The Journalled File System (JFS) was first introduced by IBM in AIX Version 3.1 to improve file system reliability. JFS features include:

- All updates to file systems are journalled to a log prior to updating the actual system
- The JFS log is replayed on reboot after a system failure and any file system updates in progress are completed
- Full file system recovery using `fsck` command is seldom necessary and thus the time to make a file system available after a failure is greatly reduced.

Logical Volume Manager

The Logical Volume Manager (LVM) subsystem manages disks at the logical level, providing:

- Management of multiple disks as a single entity
- Optimized performance
- Dynamically extendable file systems
- Software mirroring of data

7.5 History of HACMP

HACMP is designed to work as an independent control system, providing increased flexibility for configuration design and management of processor, network and application resources. HACMP has two major components:

- High Availability Subsystem

  This provides the base services for cluster membership, system management, cluster management, configuration integrity and control and base recovery services for failover and recovery.

- Concurrent Resource Manager

  This optionally adds concurrent shared access management for supported disks.

Customers may order the High Availability subsystem and Concurrent Resource Manager features separately.

The history of HACMP development is shown in Table 18.

Table 18. History of HACMP Versions

HACMP Release	Announcement Date	AIX Level Required	Mode: Concurrent or Nonconcurrent
1.1	16 June 1992	3.2.2	nonconcurrent
1.2	31 March 1993	3.2.3	nonconcurrent
		3.2.4	PTF U421400 for concurrent 7135
			PTF U425614 for concurrent 9333
2.1	21 September 1993	3.2.4	nonconcurrent
		3.2.5	PTF U426577 for concurrent
3.1	4 October 1994	3.2.5	concurrent
			no support for SSA disks
3.1.1	12 February 1995	3.2.5	concurrent
			PTF U438726 for SSA support
4.1	19 June 1995	4.1.3	concurrent for SCSI RAID disk arrays
			nonconcurrent support for serial disk subsystems
			no support for SSA disks
4.1.1	14 November 1995	4.1.4	concurrent support for SCSI RAID, serial and SSA disk subsystems

To determine the level of HACMP and PTFs installed on your system use the command:

```
lslpp -h "cluster*"
```

Note that modification levels will not necessarily be shown explicitly in the output of the `lslpp` command. It may be necessary to know that a particular PTF number corresponds to a modification level to fully identify the level of HACMP that is installed.

Before migrating, it is important to understand the differences in functionality provided between HACMP V3.1.1 and HACMP V4.1.1.

7.5.1 HACMP Version 3.1

HACMP Version 3.1 is the most recent release of HACMP running on AIX Version 3.2. Below are some HACMP features that were first introduced with HACMP V3.1.1:

- Eight-way clusters

  Eight node clusters are available in concurrent and nonconcurrent access configurations using serial disk subsystems. Four node clusters are available in concurrent and nonconcurrent access configurations using SCSI

disk subsystems. In each case, cluster configurations are dependent on bus length limitations. Prior to Version 3.1.1, clusters are limited to four nodes in all cases.

- Resource grouping

  You can logically put resources such as disks, file systems, volume groups, IP labels and applications into a single group and configure ownership and takeover on a per-group basis.

- Cascading takeover

  Cascading takeover allows a particular resource group to be backed-up by multiple nodes in an cluster. All the nodes you assign to participate in the takeover of a given resource group are part of the resource chain. Each participating node is assigned a takeover priority. When a takeover occurs, the active node with the highest priority acquires the resource group. If that node is unavailable, the node with the next-highest priority acquires the resource group and so on.

- Keepalive algorithm

  HACMP Version 3.1 contains an improved version of the keepalive algorithm that reduces traffic in large clusters.

- Conversion utilities

  Conversion utilities are provided to allow you to convert HACMP Version 1.1, 1.2 or 2.1 cluster configuration files to run under HACMP Version 3.1.

HACMP Version 3.1 requires AIX Version 3.2.5.

7.5.1.1 HACMP Version 3.1.1

HACMP Version 3.1.1 is a modification level of HACMP V3.1 and is available as a Program Temporary Fix (PTF). The main feature added by this level is support for highly available operation of RS/6000 Scalable Parallel (SP) systems. HACMP V3.1.1 also added the following enhancements:

- IP address takeover on the SP High Performance Switch
- Support for Fiber Channel Switch network adapters in RS/6000 and RS/6000 SP systems
- Support for the RS/6000 model 39H

In August 1995, further enhancements were added to HACMP V3.1.1 to support:

- Serial Storage Architecture (SSA) disks and adapters
- Enhanced SCSI-2 Fast/Wide Adapters
- SCSI target mode on Fast/Wide adapters
- RS/6000 models 591 and R21

The 3.1.1 modification level of HACMP was delivered as PTF U436331, which was the first PTF shipped for Release 3.1 of HACMP. Thus, if your HACMP system has any PTFs applied, you are implicitly running HACMP V3.1.1, however the output of the command `lslpp -L "cluster*"` will still show the level as 3.1.0.

Note

For our testing we used HACMP Version 3.1.1 as our starting HACMP level. We recommend that you upgrade your Version 3.1 cluster to the latest level before starting the migration to Version 4.1.1.

7.5.2 HACMP Version 4.1

HACMP Version 4.1 is the first release of HACMP to be supported under Version 4.1 of AIX. New features added with HACMP V4.1 include:

- Support for symmetric multiprocessing

  HACMP V4.1 supports up to eight RISC System/6000 uniprocessor or SMP systems in a highly available cluster. Support for four node concurrent access clusters using IBM SCSI disk array subsystems is also included.

- Thread-safe libraries

  All link libraries delivered with HACMP V4.1 for AIX are now thread safe allowing multiple threads in the same process to use library functions without interfering with each other.

- Support for new processors

  It supports the following new models:

  - IBM SMP server models G30, J30 and R30
  - IBM RISC System/6000 7009 model C20

- Support for new disk subsystems and adapters

  It supports the following new devices:

  - IBM 7135-210 RAIDiant Disk Array
  - IBM enhanced SCSI-2 differential fast/wide adapter

- Support for new version compatibility functions

  This function, in theory, allows you to upgrade your HACMP clusters without taking the entire cluster offline, however, as you will see in the following chapters, there are severe limitations to this support.

Like earlier versions, HACMP V4.1 also provides conversion tools to assist existing customers in converting their configuration files to run under HACMP V4.1.

The concurrent access support in HACMP V4.1 has some constraints. This release of HACMP V4.1 provides concurrent access support for the following SCSI disk array subsystems only:

- IBM 7135 RAIDiant Array Model 210
- IBM 3514 Disk Array Subsystem Models 212 or 213
- IBM 7137 Disk Array Subsystem Models 412, 413, 414, 512, 513 or 514

Each RAID device can support the attachment of up to four RISC System/6000 systems.

Support for the concurrent access on 9333 serial disk subsystem is not yet available. The new Serial Storage architecture (SSA) is not yet supported for concurrent or nonconcurrent mode.

HACMP V4.1 for AIX requires AIX 4.1.3, due to its reliance upon the base microcode levels distributed with this AIX release.

7.5.2.1 HACMP Version 4.1.1

HACMP Version 4.1.1 is a modification level of HACMP Version 4.1. This modification level adds the following new or improved features over Version 4.1.

- Concurrent access support for serial and SSA disk subsystems
- New ease-of-configuration graphical user interface
- Cluster Snapshot capability allowing the user to save and reapply complete cluster and resource configurations
- HANFS feature for providing highly available NFS server facilities
- Support for RS/6000 SP systems
- Support for RS/6000 models 591 and R21

The 4.1.1 modification level of HACMP is shipped as an installable package. HACMP must be reinstalled to migrate to the new level. All media shipped with HACMP after the GA date of HACMP 4.1.1 will include the later level. Existing customers using HACMP Version 4.1.0 can obtain the later code by ordering an MES update to their software configuration. For full details, see the announcement letter for HACMP Version 4.1.1. After upgrading, the command `lslpp -L "cluster*"` will show the installed level as 4.1.1.0.

HACMP Version 4.1.1 requires AIX Version 4.1.4 to install or run.

Attention!

We highly recommend that anyone migrating to HACMP Version 4.1 on AIX Version 4.1 ensure that they have the latest levels of both AIX and HACMP software. We experienced several problems in early testing of migration to HACMP Version 4.1.0 which were fixed in Version 4.1.1.

7.6 Granular Packaging

You will notice that HACMP Version 4.1.1, in common with most AIX Version 4.1.1 software products, is packed into a greater number of smaller filesets. This allows the system manager greater control to trade off the disk space used by the installation against the functions actually required.

The AIX Version 3.1.1 filesets and their Version 4.1.1 equivalents are given in Table 19 on page 205 below.

Table 19. HACMP Version 4.1.1 Packaging

V3.1.1 Fileset	V4.1.1 Filesets	Approx. Size (MB)	Description
cluster.client	cluster.base.client.rte	0.4	HACMP Base Client Runtime
	cluster.base.client.lib	0.5	HACMP Base Client Libraries
	cluster.base.client.utils	0.1	HACMP Base Client Utilities
	cluster.msg.en_US.client	0.1	HACMP Client Messages - U.S. English
	cluster.man.en_US.client.data	0.1	HACMP Client Man Pages - U.S. English
cluster.server	cluster.base.server.rte	2.1	HACMP Base Server Runtime
	cluster.base.server.events	1.1	HACMP Base Server Events
	cluster.base.server.utils	1.4	HACMP Base Server Utilities
	cluster.base.server.diag	0.3	HACMP Base Server Diags
	cluster.msg.en_US.server	0.2	HACMP Server Messages - U.S. English
	cluster.man.en_US.server.data	0.3	HACMP Server Man Pages - U.S. English
no equivalent	cluster.vsm.server	4.2	Visual System Management Configuration Utility
cluster.clvm	cluster.clvm.rte	0.1	HACMP for AIX Concurrent Access
	prpq.clvm	7.4	PRPQ Concurrent Logical Volume Manager
no equivalent	cluster.adt.client.demos	0.1	HACMP Client Demos
	cluster.adt.client.include	0.3	HACMP Client Include Files
	cluster.adt.client.samples.clinfo	0.4	HACMP Client CLINFO Samples
	cluster.adt.client.samples.clstat	0.1	HACMP Client Clstat Samples
	cluster.adt.client.samples.demos	0.1	HACMP Client Demos Samples
	cluster.adt.client.samples.libcl	0.3	HACMP Client LIBCL Samples
no equivalent	cluster.adt.server.demos	0.2	HACMP Server Demos
	cluster.adt.server.samples.demos	0.1	HACMP Server Sample Demos
	cluster.adt.server.samples.images	0.8	HACMP Server Sample Images

Chapter 8. HACMP Migration Process

This chapter describes upgrade considerations for migrating from HACMP Version 3.1.1 to HACMP Version 4.1.1. Particular notice should be taken of the preparation stages of documentation and determining test procedures. Care taken in planning can avoid many potential problems later in the process.

8.1 Objectives

Our goal in upgrading is to save all user configuration data whenever possible so that the HACMP cluster will behave correctly after the migration. The end user doesn't need to perform any actions to reconfigure the cluster at its new level.

Most HACMP customers require that their applications continue to run during the upgrade. In other words, at least one of the cluster nodes needs to be running at all times. So while we migrate one node, its resources are taken over by another node and are still available to users.

HACMP Version 4.1.1 makes upgrading from HACMP Version 3.1.1 less disruptive. Version compatibility, a new function in this release of HACMP, allows you to upgrade an existing cluster running HACMP Version 3.1.1 to HACMP Version 4.1.1 without taking the entire cluster offline. During the upgrade process, individual nodes in the cluster can be removed from the cluster, upgraded one at a time, then reintegrated into to the cluster. Nodes running HACMP Version 3.1.1 and HACMP Version 4.1.1 can coexist while the rest of the nodes are upgraded.

8.2 Overview

We concluded that the migration of a cluster from AIX Version 3.2.5 and HACMP Version 3.1.1 to AIX Version 4.1.4 and HACMP Version 4.1.1 is not a particularly complex procedure, however there are some points that you should consider before starting your migration:

- Ensure that you migrate your HACMP software directly to Version 4.1.1. We discovered several problems with the migration to HACMP Version 4.1.0 which were fixed in the V4.1.1 maintenance level.
- We found that it was not possible to complete the migration of your cluster without interruption to the users. Every application shared disk resource must be taken offline at some stage during the migration process.
- We strongly suggest *not* running a heterogeneous cluster of systems on mixed AIX V3.2.5 and V4.1.4 for an extended period of time due to the incompatibility of the format of the Journalled File System log volumes.

8.3 Documenting Your Cluster

It is highly recommended that your cluster environment is fully documented. It is quite likely that when the cluster was originally installed there were diagrams, tables and worksheets. It is also quite possible that your cluster has since evolved, and that the documentation is no longer correct. *Before* starting the

migration is a good time to bring your documentation back into synchronization with reality.

Much of the information required to document your cluster can be found in the HACMP smit menus accessed by using the fastpath `smit hacmp`.

```
                                    HACMP/6000

Move cursor to desired item and press Enter.

  Manage Cluster Environment
  Manage Application Servers
  Manage Node Environment
  Show Environment
  Verify Environment
  Manage Cluster Services
  Cluster Recovery Aids
  Cluster RAS Support

F1=Help             F2=Refresh          F3=Cancel           F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Much of the information is also available through AIX or HACMP commands. In order to access the HACMP commands, it is easiest if you add the directory /usr/sbin/cluster/utilities to your path.

```
PATH=$PATH:/usr/sbin/cluster/utilities
```

You may also wish to refer to the HACMP Planning Worksheets that appear in the *HACMP Planning Guide*, SC23-2700.

Some of the facets of your cluster documentation are detailed in the following sections. For samples of the outputs from the commands given in this section, see Chapter 9, “Sample HACMP Migration” on page 241.

8.3.1 Cluster Diagram

Draw a diagram showing the framework of the cluster. The purpose of the cluster diagram is to combine items such as disks and networks into one drawing that shows the cluster’s function and structure. The cluster diagram identifies the cluster name and ID, the number of nodes, IP addresses used by the nodes, the method of shared disk access and highly available applications.

8.3.2 Hostnames

You can verify the hostname of your systems by using the command `hostname`. It is always best to start with something easy!

8.3.3 Cluster ID and Name

The cluster identifier and name of your cluster can be shown by the command `cllsclstr` or through SMIT.

```
smit hacmp
  Show Environment
    Show Cluster Environment
      Show Cluster Definitions
```

8.3.4 Cluster Nodes

We can determine the cluster node names using the command `clnodename` or through the smit menu shown below.

```
smit hacmp
  Manage Cluster Environment
    Configure Nodes
      Change / Show Cluster Node Name
```

8.3.5 HACMP Adapter Configuration

We can examine the network adapters defined to HACMP by using the HACMP smit menus:

```
smit hacmp
  Manage Cluster Environment
    Configure Adapters
      Change / Show an Adapter
```

Alternatively, we can use the commands `cllsif -x` to list the defined adapters, or `cllsif` to show the details of the adapters.

8.3.6 TCP/IP Network Interfaces

To see the TCP/IP interface definitions, use `smit chinet`. Alternatively, you can use the command `lsdev -Cc if` to list the available interfaces, and then the `ifconfig` command to examine the details of each interface. For example:

```
ifconfig tr0
ifconfig tr1
ifconfig en0
```

This gives us important information on our TCP/IP configuration, such as the network masks and broadcast addresses.

Note: You will obtain different results running this command with or without HACMP running. If HACMP is running, you will see the service address on adapter tr0. If HACMP is not running, you will instead see the boot address.

8.3.7 Serial Network

Use `smit tty` to see the tty port configuration used for any HACMP serial links. You can also use the command `lsattr -El <tty name>` to obtain the same information. For example:

```
lsattr -El tty1
```

Note that the tty name is not prefixed by /dev/.

8.3.8 Shared Disk Devices

Write down the information about your shared disk technology and connections:

```
lsdev -Cc disk
```

For information on understanding the device addresses for devices attached to your system, see *Common Diagnostics Information Manual*, SA23-2765.

8.3.9 SCSI Adapter Addresses

Check the SCSI address of each SCSI adapter connected to shared disk resources using `smit chgscsi`, or the command `lsattr -El <scsi adapter name>`. For example:

```
lsattr -El scsi0
```

```
bus_mem_addr  0xe0000                  Bus memory address
dma_bus_mem   0x1200000                Address of bus memory used for DMA
dbmw          0x202000                 DMA bus memory LENGTH
bus_intr_lvl  14                       Bus interrupt level
intr_priority 3                        Interrupt priority
dma_lvl       4                        DMA arbitration level
id            7                        Adapter card SCSI ID
ucode         /etc/microcode/8d77.32.54 Name of adapter code download file
bb            no                       BATTERY backed adapter
tm            no                       Enable TARGET MODE interface
tme           no                       Target Mode interface enabled
tmp           50                       PERCENTAGE of bus memory DMA area
                                           for target mode
```

8.3.10 Shared LVM Components

Use `smit lvm` to obtain information about the physical and logical disk volumes, volume groups and file systems defined on your system. You can also use the LVM commands as shown in the examples below:

- Physical volumes

  ```
  lspv
  lspv hdisk0
  lsvg -p sharedvg
  ```

- Volume groups

  ```
  lsvg
  lsvg sharedvg
  ```

  In the output of the `lsvg sharedvg` commands, it is particularly important to note the value of the `QUORUM` and `AUTO ON` fields.

```
VOLUME GROUP:   sharedvg                VG IDENTIFIER:  00014732b5a91022
VG STATE:       active                  PP SIZE:        4 megabyte(s)
VG PERMISSION:  read/write              TOTAL PPs:      609 (2436 megaby
tes)
MAX LVs:        256                     FREE PPs:       548 (2192 megaby
tes)
LVs:            2                       USED PPs:       61 (244 megabyte
s)
OPEN LVs:       2                       QUORUM:         1
TOTAL PVs:      3                       VG DESCRIPTORS: 3
STALE PVs:      0                       STALE PPs       0
ACTIVE PVs:     3                       AUTO ON:        no
```

In addition, if you are using NFS to export any file systems in a shared volume group, you should ensure that you record the device major number of that group. Use the `ls` command as shown in the following example.

```
ls -l /dev/sharedvg
```

The output will be of the format:

```
crw-rw----   1 root     system    26,  0 Apr 05 13:42 /dev/sharedvg
```

In this case, the device major number is 26.

- Logical volumes

  ```
  lslv sharedlv
  lsvg -l sharedvg
  ```

- File systems

  ```
  lsfs
  df
  lsvg -l sharedvg
  ```

8.3.11 Resource Groups

Information about the resource chain and the individual resources that constitute each resource group can be obtained from the SMIT menus:

```
smit hacmp
  Manage Node Environment
    Manage Resource Groups
      Configure Resources for a Resource Group
```

Take note of any filesystems that are being exported using NFS.

8.3.12 Application Servers

Use SMIT to get the information about the application servers and the applications that they serve:

```
smit hacmp
  Manage Application Servers
    Change / Show an Application Server
```

8.3.13 User Defined Cluster Events

For information on customized event processing in your cluster use the SMIT menus:

```
smit hacmp
  Manage Node Environment
    Change/Show Cluster Events
```

8.3.14 Run Time Parameters

Use SMIT to examine the run time settings for HACMP on each node:

```
smit hacmp
  Manage Node Environment
    Manage Resource Groups
      Configure Run Time Parameters
```

8.3.15 Clients

Clients are end-user devices that can access the nodes in an HACMP cluster. You need to consider the following factors:

- Client programs running on PCs or other computers
- Users connecting through ASCII terminal connected to terminal servers
- Clients both with and without the Client Information Program, `clinfo`

8.4 Determining Test Procedures

Before starting the migration you should also take some time to determine how you are going to test the success or failure of the process. This is not a simple question. Although our aim is to have the cluster and applications available to users at all times, depending upon the level of redundancy and configuration of your cluster, this may not be possible. You will probably need to schedule at least some downtime to test that your event scripts are working correctly after the migration, including testing a takeover situation.

The exact testing that you will need depends on your configuration, but could include the following:

- Simulated network adapter failure by unplugging the network adapter from the network.

  If you unplug a node's service adapter, you should see a series of events occur, starting with the event *swap_adapter* as the service address is taken over by the standby adapter. The processing ends by running a *fail_standby* event, indicating there is no longer an available standby adapter in the system. When you reconnect the network adapter, you will see the event *join_standby*. This indicates the reconnected adapter has now assumed the role of standby adapter. The service address will *not* automatically swap back to the original adapter. The service address can be returned to its original adapter by either unplugging the second adapter to generate another *swap_adapter* event, or by stopping and restarting HACMP.

- Simulated network failure by introducing a break in the network. Remember if you are using a token ring, you may need to break the ring in two or more places.

  The processing in this case will depend upon the level of redundancy in your cluster. In all cases, you should see a *network_down* event. It is possible

that in the case of two-node clusters, you may instead see a *swap_adapter* event on the other node. This reflects the fact that with only two nodes it is very difficult for the Cluster Manager to distinguish between a network or an adapter failure. If there are other TCP/IP hosts on the same network, their names or addresses can be listed in the file /usr/sbin/cluster/netmon.cf. These hosts need not be part of the cluster, but will simply be sent an ICMP packet by HACMP to help determine if the complete network or only an adapter has failed.

- Complete node failure by unplugging the power from one node.

  In the case of a complete node failure, you should see other nodes in the cluster perform such actions as taking over the service address of the dead node, and accessing any shared volume groups and file systems.

- Verifying customized event scripts

  The testing to be performed for this function will focus on the processing that was added to the event scripts. There is no way for us to generalize this here.

In all testing situations be sure that you wait sufficient time to allow all resulting events to complete before starting another test.

8.5 HACMP Migration Process

In this section, we discuss the general methods used in migrating a system from AIX Version 3.2.5 and HACMP Version 3.1.1 to AIX Version 4.1.4 and HACMP Version 4.1.1.

The migrations steps are summarized below, and described in more detail in the following sections:

1. Stop HACMP on the first node

   Remove a selected node from the cluster by shutting down HACMP on the node, with takeover. This allows the other node or nodes in the cluster to takeover the resources of the chosen node, and continue serving those resources to clients while the node is upgraded.

2. Backup the selected node

   As in most major system management procedures, the first three steps are:

   a. Take a backup

   b. Check your backup

   c. Take another backup anyway

   It is important that a bootable system backup be taken before starting the migration.

3. Upgrade the operating system on the stopped node

   Because HACMP V4.1.1 requires AIX V4.1.4, you must upgrade the operating system first before installing the new level of HACMP. You will then need to perform a few manual steps before installing the new version of HACMP.

4. Upgrade HACMP on the stopped node

   In this step, you'll install HACMP V4.1.1 over the old HACMP software. Before you install the HACMP new version, you must check the disk space and pre-installed software to ensure they meet the requirements. After the

installation, you must synchronize and verify the cluster to ensure the cluster is configured correctly.

5. Restart HACMP V4.1.1 and test

   Start up HACMP V4.1.1 on the upgraded node. This node will then re-gain its resources. You can now test HACMP V3.1.1 and HACMP V4.1.1 interoperability for node failure, network adapter failure and network failure. Depending on the amount of testing of applications and system management procedures that you were able to perform before starting the migration, you may wish to pause the migration at this point to allow time for a full load testing of your application running under HACMP V4.1.1 before committing all nodes to the new software level. However, for reasons that will be explained later, we do not recommend running in a heterogeneous environment for long periods of time.

6. You can then repeat this procedure on the next node in your cluster

8.5.1 Stopping HACMP on the First Node

Use SMIT to stop HACMP.

```
smit hacmp
  Manage Cluster Services
    Stop Cluster Services
```

You will see the following screen.

```
                              Stop Cluster Services

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

                                                      [Entry Fields]
* Stop now, on system restart or both                 now                      +

  BROADCAST cluster shutdown?                         true                     +
* Shutdown mode                                       graceful                 +
          (graceful, graceful with takeover, forced)

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

Be sure the Shutdown mode field is set to takeover. This will allow the other node or nodes in the cluster to assume the resources of the selected node, allowing applications and other resources to remain available to the user during the migration process.

Alternatively, you can stop HACMP from the command line:

```
/usr/sbin/cluster/utilities/clstop -y -N -gr
```

8.5.2 Backing Up the Node

In the event of a failure in the migration process, a good backup is vital to the fast resumption of full services to your users. You should take a bootable system backup following the procedures detailed in the redbook *A Holistic Approach to AIX V4.1 Migration, Volume 1*, SG24-4652.

Note

It is important that you take your system backup while HACMP is *not* running. This will ensure that the node's boot address is active, and will be used at installation and reboot time if the system is reinstalled from the system backup.

Although data on non-rootvg volume groups should not be affected by the migration, we *strongly recommend* that you ensure you have a reliable backup of this information before starting the migration process. For data on shared volume groups, this backup must be made before stopping HACMP in order to have the volume groups accessible to the system.

8.5.3 Upgrading AIX

Use the AIX Migration installation method to upgrade the operating system to AIX V4.1.4. For information on the AIX upgrade process, see the redbook *A Holistic Approach to AIX V4.1 Migration, Volume 1*, SG24-4652. After the migration, you will have to manually take some action to correct problems caused by the combination of the upgrade process and HACMP.

8.5.3.1 Problems Encountered After the AIX Upgrade

When migrating a system running HACMP from AIX Version 3.2 to AIX Version 4.1.4, you will probably encounter the problems described in the following sections. In each case, the problem is not severe as long as you are aware of it and can be prepared to take the necessary action.

Volumes on Shared Disk: In most cases of clusters running HACMP, you will have at least one volume group configured as a takeover resource. In this case, two or more systems will have knowledge of the volume group and the file systems that it contains. Each system is set to *not* vary on the volume group at boot time. HACMP then has control of the volume group and will vary it on to the owning node in normal operations. In a node failure situation, HACMP will vary the volume group on to the backup node, or the node with the next highest priority in the cluster.

Part of the migration installation process involves the recreation of the volume group and file system information in the Object Data Manager (ODM). To achieve this, non-rootvg volume groups are imported at the end of the migration process. This reads the information about the volume group from the Volume Group Descriptor Area (VGDA), which is found on each disk in the volume group, and stores that information in the ODM.

In the case of a migration installation of a node in an HACMP cluster, however, the non-rootvg volume groups will be varied on to the backup node that is providing service to the users. In this case the import of the volume group fails because the backup system holds a reserve condition on the disks in the volume group.

To regain the configuration of the non-rootvg volume groups on the migrated system it is necessary to manually perform the following steps for each non-rootvg volume group.

Note

This procedure requires that the file systems in the non-rootvg volume group, and thus the applications that use the data in these file systems, be taken offline. That is, continuous availability to users through the entire migration process is *not* possible.

On the backup node that is currently accessing the non-rootvg volume group:

1. Ensure that all users of data on the shared volume group are off the system, and that applications using this data are stopped.
2. Unmount the file systems in the shared volume group.
3. Vary off the volume group using the command `varyoffvg vg_name`.

On the migrated node:

4. Import the volume group.

   The volume group can be imported using SMIT:

   `smit importvg`

   SMIT returns the following screen.

```
                              Import a Volume Group

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

                                                      [Entry Fields]
  VOLUME GROUP name                                  [sharedvg]
* PHYSICAL VOLUME name                               [hdisk3]                +
* ACTIVATE volume group after it is                   yes                    +
    imported?
  Volume Group MAJOR NUMBER                          [33]                    +#

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

   Enter the field values as follows:

   VOLUME GROUP name
   Enter the name of the volume group that you are importing.

   PHYSICAL VOLUME name
   Enter the name of one of the physical volumes that resides in the volume group. Refer to the list of physical volumes you recorded when documenting your volume group.

ACTIVATE volume group after it is imported?

Set this field to **yes**

Volume Group MAJOR NUMBER

If you are not using NFS to export any of the file systems in this volume group, you can leave this field blank to accept the default value. If you are exporting some of the file systems, or think you may require this function in the future, see “Device Major Numbers” on page 218 below to determine the correct setting for this field.

Press Enter to commit the information.

5. Set the volume group to match previous settings.

   By default, a volume group that has just been imported is configured to automatically become active at system restart. In an HACMP environment, a volume group should be varied on as appropriate by the cluster event scripts.

   In addition, the default setting is to ensure that a quorum of volume group descriptor areas (defined as *greater than 50%*) is available before allowing the volume group to be varied on. This setting is optional and is often changed in HACMP environments.

   Therefore, after importing a volume group, use the SMIT menu to reconfigure the volume group to match its state before the migration. You should have recorded these values in 8.3.10, “Shared LVM Components” on page 210.

   ```
   smit chvg
   ```

   SMIT prompts you to identify the volume group. Press F4 to list the defined volume groups. Select the appropriate volume group from the list and press Enter.

   The following screen appears. The first item is filled in with the volume group name you specified in the previous screen.

```
                          Change a Volume Group

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

                                                      [Entry Fields]
* VOLUME GROUP name                                   sharedvg
* Activate volume group AUTOMATICALLY                 no                      +
    at system restart?
* A QUORUM of disks required to keep the volume       yes                     +
   group on-line ?

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

   Refer to the documentation collected in 8.3.10, “Shared LVM Components” on page 210 for the correct values for these fields.

Press Enter to commit the values.

Note: Remember that these values will be reset every time the volume group is exported and reimported. They will not be lost if the volume group is simply varied off and varied on.

6. Now, on the migrated node, vary off the volume group using the command `varyoffvg <vg name>`.

Then, on the backup node:

7. Vary on the volume group using the command `varyonvg <vg name>`.

Once you have completed the above steps for each shared volume group, you can resume services to your users from the backup node while you complete the other steps in the migration.

Device Major Numbers: If you are using NFS to export the file systems in a shared volume group, you need to ensure that the volume group uses the same device major number on each node in the cluster. The major number is used as part of the NFS file handle. The important factor here is that if the device major numbers on the owning system and the takeover system are the same, client systems using NFS to mount file systems in the volume group will not lose their connections in the event of a node failure and takeover—they will see only a delay, then their file accesses will complete. If the major numbers differ, clients will have to unmount then re-mount the file system to regain access after a takeover.

In 8.3.10, “Shared LVM Components” on page 210, you should have recorded the major number previously used for the volume group. You can verify that this major number is available on the migrated system by using the command `lvlstmajor`. An example of the output from an AIX V4.1.4 system is shown below.

```
33...
```

In this sample output, major numbers of 33 or above are available for use.

The problem arises from the fact that the default assignments of device major numbers have changed between AIX V3.2.5 and V4.1.4. It is quite likely the major number used previously in your cluster will not be available on the migrated system. In this case you will have to select a new device major number. Use the `lvlstmajor` command on the migrated system to select a suitable value. We recommend that you leave a few vacant numbers before your selection to allow for the case where other systems in the cluster have more devices and thus use more of the acceptable range for other purposes.

Once you have determined you must change your device major number, a new problem arises. If you are migrating all nodes at the same time, you can simply use the new value as you upgrade each node. However, if you intend running heterogeneously for any period of time, you must consider the fact that using different major numbers does add an additional step to a takeover situation - that of unmounting and re-mounting NFS file systems on each of your clients.

To avoid having to perform this additional step in a crisis situation, you can change the device major number on the AIX V3.2.5 and HACMP V3.1.1 nodes. This must be performed while you have the users off the cluster, but can be scheduled for a convenient time. Refer to the procedure documented in

"Volumes on Shared Disk" on page 215 above. After step 6, varying off the volume group, you should export and reimport the volume group on any other nodes in the cluster to which the volume group is defined. When reimporting the volume group, specify the newly chosen major number. Don't forget to also change the volume group settings on each system after importing.

SCSI Adapter Addresses: Every device connected to a particular SCSI bus, including the SCSI adapter or adapters, must have a unique SCSI address. The standard address assigned to the SCSI adapter is 7.

In an HACMP environment with shared SCSI disks, at least one of the connected adapters (there could be two or more) must be set to a different address. The names of the SCSI adapters on the system can be seen in the output of the command:

```
lsdev -Cc adapter
```

The SCSI address assigned to a particular adapter can be shown with the command:

```
lsattr -El <adapter name> -a id
```

For example:

```
# lsdev -Cc adapter
sio0   Available 00-00    Standard I/O Planar
scsi0  Available 00-02    SCSI I/O Controller
scsi1  Available 00-06    SCSI I/O Controller
scsi2  Available 00-07    SCSI I/O Controller
scsi3  Available 00-08    SCSI I/O Controller
fda0   Available 00-00-0D Standard I/O Diskette Adapter
siokb0 Available 00-00-0K Keyboard Adapter
siotb0 Available 00-00-0T Tablet Adapter
sa0    Available 00-00-S1 Standard I/O Serial Port 1
sa1    Available 00-00-S2 Standard I/O Serial Port 2
gda0   Available 00-01    Color Graphics Display Adapter
tok0   Available 00-03    Token-Ring High-Performance Adapter
tok1   Available 00-04    Token-Ring High-Performance Adapter
slc0   Available 00-00    Serial Optical Link Chip
sioms0 Available 00-00-0M Mouse Adapter
ppa0   Available 00-00-0P Standard I/O Parallel Port Adapter

# lsattr -El scsi0 -a id
id 7 Adapter card SCSI ID True

# lsattr -El scsi1 -a id
id 7 Adapter card SCSI ID True

# lsattr -El scsi2 -a id
id 6 Adapter card SCSI ID True

# lsattr -El scsi3 -a id
id 7 Adapter card SCSI ID True
```

When the system is reinstalled, the address of the SCSI adapter will revert to the default address for that type of SCSI adapter. If there is already another adapter on the same SCSI bus that is using the same address, there will be a conflict. This problem also occurs if the system is booted from diskette, CD-ROM or tape for performing system maintenance. For this reason, it is strongly recommended that where possible, no devices on the SCSI bus should have a permanent

address of 7. In this case, an adapter being serviced can happily use the default value until being set to its permanent value, without causing any conflicts.

If one of the adapters connected to a shared SCSI bus in your cluster uses the default address, you should assign it a new address. We recommend that you choose this node to upgrade first. That way, the default value will not cause any problems. During the course of the upgrade, before importing the volume group from the disks, you should change the SCSI address to the new assigned value. Check to see if the devices attached to the adapter are in the **Defined** or **Available** states by using the command `lsdev -Cs scsi`:

```
rmt0   Available 00-08-00-30 2.3 GB 8mm Tape Drive
hdisk0 Available 00-08-00-00 670 MB SCSI Disk Drive
hdisk1 Available 00-08-00-40 355 MB SCSI Disk Drive
cd0    Available 00-08-00-50 CD-ROM Drive
hdisk2 Available 00-06-00-00 1.0 GB Differential SCSI Disk Drive
hdisk3 Available 00-06-00-10 1.0 GB Differential SCSI Disk Drive
hdisk4 Available 00-07-00-00 857 MB SCSI Disk Drive
hdisk5 Available 00-07-00-10 857 MB SCSI Disk Drive
hdisk6 Available 00-07-00-20 857 MB SCSI Disk Drive
```

You can change the address one of two ways. The first method is to use the command:

```
chdev -l <adapter name> -a id=<new address> -P
```

or the

```
smit chgscsi
```

panel:

```
                Change / Show Characteristics of a SCSI Adapter

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

                                                        [Entry Fields]
  SCSI Adapter                                         scsi1
  Description                                          SCSI I/O Controller
  Status                                               Available
  Location                                             00-06
  Adapter card SCSI ID                                [6]                      +#
  BATTERY backed adapter                               no                      +
  DMA bus memory LENGTH                               [0x202000]               +
  Enable TARGET MODE interface                         no                      +
  Target Mode interface enabled                        no
  PERCENTAGE of bus memory DMA area for target mode   [50]                     +#
  Name of adapter code download file                   /etc/microcode/8d77.a0>
  Apply change to DATABASE only                        yes                     +

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

Ensure that you either use the -P option on the command above, or select **yes** for the **Apply change to DATABASE only** field. This will not attempt to change the running system, but will change the ODM database. You must then reboot the system to use the new value.

If you wish to perform the change without rebooting the system, you must use the following procedure.

1. If any of the devices on that SCSI bus are Available, they must be changed to the defined state by using the command rmdev -l <device name>. For example,

```
# rmdev -l hdisk2
hdisk2 Defined

# rmdev -l hdisk3
hdisk3 Defined

# lsdev -Cs scsi
rmt0   Available 00-08-00-30 2.3 GB 8mm Tape Drive
hdisk0 Available 00-08-00-00 670 MB SCSI Disk Drive
hdisk1 Available 00-08-00-40 355 MB SCSI Disk Drive
cd0    Available 00-08-00-50 CD-ROM Drive
hdisk2 Defined   00-06-00-00 1.0 GB Differential SCSI Disk Drive
hdisk3 Defined   00-06-00-10 1.0 GB Differential SCSI Disk Drive
hdisk4 Available 00-07-00-00 857 MB SCSI Disk Drive
hdisk5 Available 00-07-00-10 857 MB SCSI Disk Drive
hdisk6 Available 00-07-00-20 857 MB SCSI Disk Drive
```

2. You can now change the SCSI address of the adapter using the command

   chdev -l <adapter name> -a id=<new address>

or by using the

`smit chgscsi`

panel:

```
                Change / Show Characteristics of a SCSI Adapter

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

                                                      [Entry Fields]
  SCSI Adapter                                       scsi1
  Description                                        SCSI I/O Controller
  Status                                             Available
  Location                                           00-06
  Adapter card SCSI ID                              [6]                          +#
  BATTERY backed adapter                             no                          +
  DMA bus memory LENGTH                             [0x202000]                   +
  Enable TARGET MODE interface                       no                          +
  Target Mode interface enabled                      no
  PERCENTAGE of bus memory DMA area for target mode [50]                         +#
  Name of adapter code download file                 /etc/microcode/8d77.a0>
  Apply change to DATABASE only                      no                          +









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

3. After changing the adapter you can now re-activate the disks using the command `mkdev -l <disk name>`. For example,

```
# mkdev -l hdisk2
hdisk2 Available

# mkdev -l hdisk3
hdisk3 Available

# lsdev -Cs scsi
rmt0   Available 00-08-00-30 2.3 GB 8mm Tape Drive
hdisk0 Available 00-08-00-00 670 MB SCSI Disk Drive
hdisk1 Available 00-08-00-40 355 MB SCSI Disk Drive
cd0    Available 00-08-00-50 CD-ROM Drive
hdisk2 Available 00-06-00-00 1.0 GB Differential SCSI Disk Drive
hdisk3 Available 00-06-00-10 1.0 GB Differential SCSI Disk Drive
hdisk4 Available 00-07-00-00 857 MB SCSI Disk Drive
hdisk5 Available 00-07-00-10 857 MB SCSI Disk Drive
hdisk6 Available 00-07-00-20 857 MB SCSI Disk Drive
```

For information on understanding the device addresses for SCSI devices and adapters attached to your system, see *Common Diagnostics Information Manual*, SA23-2765.

Target Mode SCSI Enablement: Target mode SCSI is a facility that allows keepalive packets to be transferred between nodes across a shared SCSI bus.

To enable target mode SCSI operations, a setting must be made under the SCSI adapter smit panel:

```
smit
  Devices
    SCSI Adapter
      Change / Show Characteristics of a SCSI Adapter
```

Select the appropriate SCSI adapter and press Enter.

```
                Change / Show Characteristics of a SCSI Adapter

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

                                                      [Entry Fields]
  SCSI Adapter                                       scsi1
  Description                                        SCSI I/O Controller
  Status                                             Available
  Location                                           00-06
  Adapter card SCSI ID                              [5]                      +#
  BATTERY backed adapter                             no                      +
  DMA bus memory LENGTH                             [0x202000]               +
  Enable TARGET MODE interface                       yes                     +
  Target Mode interface enabled                      yes
  PERCENTAGE of bus memory DMA area for target mode [50]                     +#
  Name of adapter code download file                 /etc/microcode/8d77.a0>
  Apply change to DATABASE only                      no                      +



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

This setting will be lost during the AIX upgrade, and will revert to the default value of no. If you use SCSI target mode on your nodes, you must set this value by hand.

The change must either be made in the database only, and the system rebooted, or the adapter can be changed online using a procedure like the one described in “SCSI Adapter Addresses” on page 219.

TCP/IP Setup: This section describes how HACMP changes the file /etc/inittab and /etc/rc.net, and how this affects the starting of TCP/IP after the AIX migration.

The /etc/inittab File: When IP address takeover is configured by defining a service adapter as a resource, the HACMP system edits /etc/inittab to change the rc.tcpip and inet dependent entries from runlevel ″2″ (the default multi-user level) to runlevel ″a″. Entries that have level ″a″ are processed only when the telinit command requests them to be run. In the /etc/inittab file the following entries are changed or modified by HACMP.

```
rtcpip:a wait:/etc/rc.tcpip > /dev/console 2>&1 # Start TCP/IP daemons
rcnfs:a:wait:/etc/rc.nfs > /dev/console 2>&1 # Start NFS Daemons
rcncs:a:wait:sh /etc/rc.ncs
qdaemon:a:wait:/usr/bin/startsrc -sqdaemon
writesrv:a:wait:/bin/startsrc -swritesrv
harc:2:wait:/usr/sbin/cluster/etc/harc.net #HACMP for AIX network startup
clinit:a:wait:touch /usr/sbin/cluster/.telinit # HACMP for AIX This must
be last entry in inittab!
```

After the migration from AIX V3.2.5 to AIX V4.1.4, the operating system creates a new /etc/inittab file without the above HACMP related entries and restores the settings to runlevel ″2″. For example:

```
rtcpip:2 wait:/etc/rc.tcpip > /dev/console 2>&1 # Start TCP/IP daemons
rcnfs:2:wait:/etc/rc.nfs > /dev/console 2>&1 # Start NFS Daemons
qdaemon:2:wait:/usr/bin/startsrc -sqdaemon
writesrv:2:wait:/bin/startsrc -swritesrv
```

The /etc/rc.net File: HACMP edits the /etc/rc.net file on each node for which IP address takeover might occur. The /etc/rc.net file is modified so that it can only be executed by the Cluster Manager when starting HACMP for AIX, instead of being executed during the AIX boot process.

The system places the following code at the start of the /etc/rc.net script.

```
# HACMP for AIX
# HACMP6000
# HACMP6000 These lines added by HACMP6000 software
 [ ″$1″ = ″-boot″ ] && shift || exit 0 # HACMP6000
# HACMP6000
```

This code causes the script to exit unless it was called with an additional -boot parameter. This will not be present unless the script is called by HACMP.

In HACMP, the /etc/rc.net file is called by the /usr/sbin/cluster/etc/rc.cluster file to configure the network. If /usr/sbin/cluster/etc/rc.cluster is called with the -boot parameter and a boot address is defined for a node, then /usr/sbin/cluster/etc/rc.cluster calls /etc/rc.net with the -boot parameter, and the configuration of the network will complete. This is how the HACMP entry in the /etc/inittab file is defined to work.

After Migration: After AIX migration from V3.2.5 to V4.1.4, the HACMP modifications to the /etc/rc.net file are still present, however the HACMP modifications to the /etc/inittab file have been removed as the inittab file was replaced during the migration. Thus the /etc/inittab file and /etc/rc.net file are not consistent. In this case, if you reboot the new AIX system, the /etc/rc.net file is not called by HACMP Cluster Manager with the -boot option. It is called by the AIX boot process without the -boot option. The script will immediately exit and the TCP/IP networks will not be started. You will see error messages from TCP/IP, such as:

```
Sendmail daemon: sendmail: 0832-035 Cannot create a socket: No protocol
of the specified type & domain exists.
```

To solve this problem, you should remove the HACMP lines from the /etc/rc.net file. This is a temporary state—the lines will be replaced automatically later in the migration process, when HACMP is reconfigured. Removing these lines allows your network connections to become active. The network connections will be required later in the process to allow you to synchronize the configuration of the node with other nodes in the cluster.

Notes

On the first reboot after migrating your system to AIX Version 4.1.4 (when you are presented with the license information and installation assistant), you will find that TCP/IP is running correctly and this problem will not exhibit. This is because the first boot follows a different boot process than a regular boot. The problem will still occur on subsequent reboots unless you remove the lines as recommended above.

If you proceed to install the new HACMP code at this stage, the HACMP installation will itself remove these lines. They will not be re-inserted until HACMP is reconfigured to include IP address takeover.

When you have finished the work mentioned above, the operating system is ready for the HACMP upgrade.

8.5.4 Upgrading HACMP

This section provides instructions for upgrading an HACMP for AIX configuration by overwriting the previous version of the software and resynchronizing your configuration. The process can be broken into several parts:

1. Prerequisites

   This includes checking for sufficient disk space, as well as installing BOS filesets required by HACMP.

2. Preparing for the Upgrade

   In this step you prepare for the upgrade procedure by saving your configuration, and by ensuring that your current installation is in the committed state.

3. Installing HACMP for AIX Software

   Again, there are some steps that must be taken after the migration to ensure that you do not have problems synchronizing your cluster and running the cluster in the future.

4. Synchronizing the Cluster

   For the HACMP cluster to function properly, the HACMP ODM entries must be the same on all cluster nodes. You must synchronize all cluster nodes whenever you make a change to the cluster′s definition. If the definitions are not synchronized across nodes, the HACMP daemons will not start on the cluster nodes.

5. Recover the ODM Event Database

   The cluster event configurations are not preserved in the migration process, and are not copied from the other nodes in a synchronize operation. Thus you must restore your ODM event database or manually reconfigure your customized event scripts.

6. Verifying the Environment

   After reconfiguring a cluster or updating a node environment, run the Cluster Verification procedure on one node to check that all resources used by HACMP are validly configured and that ownership and takeover of those resources are defined and agreed upon by all nodes.

8.5.4.1 Prerequisites

The following prerequisites must be satisfied before installing HACMP.

- The /usr directory must have 3 MB of free space for nodes in a non-concurrent environment and 6 MB for nodes in a concurrent access environment. If you are only installing the HACMP for AIX client software, 2 MB are required. These figures are for the additional disk space required above that used by the V3.1.1 level. The total amount of space needed for HACMP for AIX Version 4.1.1 is approximately 15 MB , 18 MB, and 8 MB for non-concurrent, concurrent and client environments respectively.
- Among other prerequisites, HACMP V4.1.1 requires that the fileset bos.compat.lan be installed. In most cases, the bos.compat.lan fileset will not be installed automatically during an AIX migration installation. HACMP also requires the following compatibiility filesets:
  - bos.compat.cmds
  - bos.compat.libs
  - bos.compat.links
  - bos.compat.net

  However, these filesets will be automatically installed by the AIX migration process. After AIX migration, type the `lslpp -h "bos.compat*"` command to check if the bos.compat.lan fileset is installed. If not, install the fileset from your AIX installation media by following the steps below. Because the installation of bos.compat.lan changes the operating system kernel, you must reboot the system after installation.

  1. Insert your AIX V4.1.4 installation media and enter the command `smit install_selectable_all`. Select the name of your installation device.

```
                Install New Software Products at Latest Level

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

                                                      [Entry Fields]
* INPUT device / directory for software               /dev/rmt0.1
* SOFTWARE to install                                 [all_licensed]          +
  PREVIEW only? (install operation will NOT occur)    no                      +
  COMMIT software updates?                            yes                     +
  SAVE replaced files?                                no                      +
  ALTERNATE save directory                            []
  AUTOMATICALLY install requisite software?           yes                     +
  EXTEND file systems if space needed?                yes                     +
  OVERWRITE same or newer versions?                   no                      +
  VERIFY install and check file sizes?                no                      +
  Include corresponding LANGUAGE filesets?            yes                     +
  DETAILED output?                                    no                      +



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

2. To fill in the **SOFTWARE to install** field, press F4. SMIT searches the installation media and lists all of the filesets available. Use the find function, by pressing the / key, to locate the *LAN COMIO Compatibility Software* entry, then use the F7 key to select that fileset.
3. Enter values for the other fields as shown above.
4. When you are satisfied with the entries, press Enter to install the fileset.
5. As mentioned in the install messages, after installing the bos.compat.lan fileset, you should reboot your RISC System. Ensure that you have read “TCP/IP Setup” on page 223 before you attempt to reboot.

8.5.4.2 Preparations for Upgrading

Complete the following steps before upgrading to HACMP Version 4.1.1 for AIX.

1. For each node, archive the /usr/sbin/cluster directory to a readily accessible place on disk. This allows you to easily retrieve and compare localized script and configuration files. This can be done by simply copying the directory, or by using the `tar`, `backup`, `cpio` or `pax` commands.
2. When migrating your nodes, any customized events that you have defined will be lost. To prevent this, if you use customized events in your cluster, you should save the event ODM database using the command:

   ```
   cp /etc/objrepos/HACMPevent /etc/objrepos/HACMPevent.save
   ```

   :i2refid=ha.event, saving customized scripts
3. If the installation is applied but not committed, commit it so that Version 4.1.1 can be installed over the existing version. To see if your HACMP software is already committed, enter the command:

   ```
   lslpp -L "cluster*"
   ```

   The output will be of the form:

```
  Fileset                      Level  State  Description
  ----------------------------------------------------------------------------
  cluster.client             3.1.0.0    C    HACMP/6000
                     3.1.0.1.U439890    A    HACMP/6000
  cluster.clvm               3.1.0.0    C    HACMP/6000
  cluster.server             3.1.0.0    C    HACMP/6000
                     3.1.0.1.U439890    A    HACMP/6000


State Codes:
 A -- Applied.
 B -- Broken.
 C -- Committed.
 O -- Obsolete.  (partially migrated to newer version)
 ? -- Inconsistent State...Run lppchk -v.
```

If the software and all updates are not committed, run the `smit install_commit` command before installing the Version 4.1.1 software. You will see the following screen:

```
              Commit Applied Software Updates (Remove Saved Files)

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

                                                      [Entry Fields]
* SOFTWARE name                                      [all]                  +
  PREVIEW only? (commit operation will NOT occur)     no                     +
  COMMIT requisites?                                  yes                    +
  DETAILED output?                                    no                     +









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

Press F4 to see the list of software that is in the applied state. Use the arrow keys or search for the string *cluster* to locate HACMP. Use F7 to select all HACMP components (those starting with *cluster*). Press Enter to accept the selections, then Enter again to process the commit.

8.5.4.3 Installing HACMP V4.1.1 on the First Node

Complete the following steps to install High Availability software:

1. Insert the HACMP for AIX media and enter the command `smit install_selectable_all`. Select your installation device or directory and press Enter.

```
                Install New Software Products at Latest Level

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

                                                      [Entry Fields]
* INPUT device / directory for software               /dev/rmt0.1
* SOFTWARE to install                                 [all_licensed]          +
  PREVIEW only? (install operation will NOT occur)    no                      +
  COMMIT software updates?                            yes                     +
  SAVE replaced files?                                no                      +
  ALTERNATE save directory                            []
  AUTOMATICALLY install requisite software?           yes                     +
  EXTEND file systems if space needed?                yes                     +
  OVERWRITE same or newer versions?                   no                      +
  VERIFY install and check file sizes?                no                      +
  Include corresponding LANGUAGE filesets?            yes                     +
  DETAILED output?                                    no                      +



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

2. Select the software to install. If HACMP is on the media alone, you can use the default value of **all_licensed**. Otherwise, press F4 to view a list of available software and choose the HACMP filesets by pressing F7 beside each package, or beside the HACMP all line. Press Enter to accept the selections. Enter values for the other fields as shown above. Press Enter to perform the installation.

After installing the HACMP software you should reboot the system to activate the global ODM.

8.5.4.4 Upgrading HACMP V3.1.1 to HACMP V4.1.1

During the HACMP upgrade, several steps are taken to preserve the existing configuration, and to install and configure the new HACMP code. These steps are by the install process running the pre-removal, pre-installation, post-installation and configuration scripts for the various HACMP filesets. For a description of the Licensed Program Product installation process, and the functions of these scripts, see the redbook *A Holistic Approach to AIX V4.1 Migration, Volume 1*, SG24-4652 The most important steps taken in the installation of the cluster.base filesets, and the name of the scripts that perform these steps, are listed below. You will see these steps mentioned in the output from the installation command.

`cluster.base.server.rte.pre_rm`

1. Converting the HACMP ODM databases into OLD ones

   This process involves both copying the databases to new names as shown below, and modifying the data inside the database, if required, to be consistent with the new filename.

```
Converting previously created HACMP ODM object classes to OLD ones
Converting HACMPadapter ODM class to HACMPadapterOLD
Converting HACMPcluster ODM class to HACMPclusterOLD
Converting HACMPcommand ODM class to HACMPcommandOLD
Converting HACMPevent ODM class to HACMPeventOLD
Converting HACMPfence ODM class to HACMPfenceOLD
Converting HACMPgroup ODM class to HACMPgroupOLD
Converting HACMPnetwork ODM class to HACMPnetworkOLD
Converting HACMPnim ODM class to HACMPnimOLD
Converting HACMPnode ODM class to HACMPnodeOLD
Converting HACMPresource ODM class to HACMPresourceOLD
Converting HACMPserver ODM class to HACMPserverOLD
Converting HACMPsp2 ODM class to HACMPsp2OLD
```

2. Saving the configuration files

   The installation process will now save the files defined by each new fileset as files from the previous version that are considered configuration files. The following list shows the configuration files that will be saved before the installation of the listed fileset.

```
cluster.base.client.rte:

/etc/rc.cluster                             user_merge
/usr/sbin/cluster/etc/rc.cluster            user_merge
/usr/sbin/cluster/clhosts                   user_merge
/usr/sbin/cluster/etc/clhosts               user_merge
/usr/sbin/cluster/samples/clinfo.rc         user_merge
/usr/sbin/cluster/etc/clinfo.rc             user_merge

cluster.base.client.utils:

/usr/sbin/cluster/clexit.rc                   user_merge
/usr/sbin/cluster/utilities/clexit.rc         user_merge

cluster.base.server.events:

[ all of the event scripts from the /usr/sbin/cluster/events and
/usr/sbin/cluster/samples directories ]     user_merge

cluster.base.server.utils:

/usr/sbin/cluster/utilities/clchipat        user_merge
/usr/sbin/cluster/clchipat                  user_merge
```

   Note that each of these configuration files are marked for *user_merge* processing. This means that the old files will be saved in their usual directory structures under the directory /usr/lpp/save.config, however they will not be processed in any way. If the user has made some site specific changes to these files, they must be examined and made again to the new configuration files.

`cluster.base.client.rte.pre_i`

3. Creating the cluster.log file

   If the file /usr/adm/cluster.log does not exist, it will be created.

`cluster.base.server.utils.config`

4. Adding a cron entry

   After deleting any cron entries for `clcycle` and `cllvm`, the following entry will be added to the root crontab:

```
0 0 * * * $CLUSTDIR/utilities/clcycle 1>/dev/null 2>/dev/null \
# HACMP for AIX Logfile rotation
```

cluster.base.server.rte.post_i

5. Removing old SMIT panels

   The ODM entries for the HACMP smit panels will be removed. At this time, the trace entries for HACMP are also defined.

cluster.base.client.rte.config

6. Adding entries to /etc/services

   As this is processing for the client fileset, only the entry for clinfo_deadman will be made at this point.

7. Adding clinfo subsystem and notify method

   The clinfo subsystem will be defined to the system resource controller, including the clinfo notify method.

cluster.base.server.rte.config

8. Adding entries to /etc/services

   For the server fileset, the following entries are added:

```
clm_keepalive   6000/udp
cllockd         6100/udp
clm_pts         6200/tcp
clsmuxpd        6270/tcp
clm_lkm         6150/tcp
clm_smux        6175/tcp
godm            6177/tcp
```

9. Adding Global ODM (GODM) to /etc/inetd.conf

   At this time, the godm is added as a sub server under the inetd super server. Note, this is added regardless of the fact that the entry may already be present. A duplicate entry here will cause problems synchronizing and verifying your configuration. For more information on this problem, see 8.5.4.5, “Problems Encountered After the HACMP Upgrade” on page 232.

10. Adding HACMP entries to /etc/snmp.conf and /etc/snmp.peers

    These entries allow HACMP to pass information to a monitoring program through the Simple Network Management Protocol (SNMP).

11. Adding loopback entry to clhosts file

    An entry for the default TCP/IP loopback interface (127.0.0.1) is added to the file /usr/sbin/cluster/etc/clhosts.

12. Adding entries to tcp.clean

    Entries are then added to the script /etc/tcp.clean to ensure that HACMP will be stopped together with TCP/IP. In most cases, this will occur if the system is shut down without first stopping HACMP.

13. Defining HACMP server subsystems

    The following subsystems are defined to the SRC:

```
0513-071 The clstrmgr Subsystem has been added.
0513-068 The clstrmgr Notify method has been added.
0513-071 The cllockd Subsystem has been added.
0513-068 The cllockd Notify method has been added.
0513-071 The clsmuxpd Subsystem has been added.
0513-068 The clsmuxpd Notify method has been added.
```

14. Removing entries from inittab and rc.net

    Any existing entries are removed from these files. These entries will be recreated during the configuration of HACMP if IP address takeover is used. For more information, see “TCP/IP Setup” on page 223.

15. Adding HACMP entry to the syslog.conf file

8.5.4.5 Problems Encountered After the HACMP Upgrade

The configuration data for all nodes in the HACMP cluster is maintained in Object Data Manager (ODM) databases on each node. The Global Object Data Manager (GODM) is the facility used by HACMP to keep these nodes synchronized. It allows the nodes to be updated through TCP/IP over a network.

The GODM runs under the control of the internet super daemon inetd When HACMP is installed, a line is added to the inetd configuration file /etc/inetd.conf to allow the super daemon to invoke the GODM when required, as shown in the following extract of the /etc/inetd.conf file:

```
chargen dgram   udp wait    root    internal
daytime dgram   udp wait    root    internal
time    dgram   udp wait    root    internal
## The following line is for installing over the network.
#instsrv stream tcp nowait  netinst /u/netinst/bin/instsrv instsrv -r /tmp/netin
stalllog /u/netinst/scripts
executiond  sunrpc_tcp tcp  wait    root    /usr/lpp/sd/executiond executiond 30
0201 1
comp_ed sunrpc_tcp tcp  wait    root    /usr/lpp/sd/executiond comp_ed 33333332
1
godm    stream  tcp nowait  root    /usr/sbin/cluster/godmd
ttdbserver  sunrpc_tcp  tcp wait    root    /usr/dt/bin/rpc.ttdbserver rpc.ttdbs
erver 100083 1
```

When migrating to HACMP Version 4.1.1, this line will be added again to the file, leaving two lines. This causes problems when attempting to synchronize and verify the cluster configuration.

```
chargen dgram   udp wait    root    internal
daytime dgram   udp wait    root    internal
time    dgram   udp wait    root    internal
## The following line is for installing over the network.
#instsrv stream tcp nowait  netinst /u/netinst/bin/instsrv instsrv -r /tmp/netin
stalllog /u/netinst/scripts
executiond  sunrpc_tcp tcp  wait    root    /usr/lpp/sd/executiond executiond 30
0201 1
comp_ed sunrpc_tcp tcp  wait    root    /usr/lpp/sd/executiond comp_ed 33333332
1
godm    stream  tcp nowait  root    /usr/sbin/cluster/godmd
ttdbserver  sunrpc_tcp  tcp wait    root    /usr/dt/bin/rpc.ttdbserver rpc.ttdbs
erver 100083 1
godm    stream  tcp nowait  root    /usr/sbin/cluster/godmd
```

Attention

Ensure that you remove (or comment out using the # symbol) the duplicate GODM entry from the /etc/inetd.conf after installing HACMP V4.1.1, but before attempting any configuration.

After removing the entry you should signal `inetd` to re-read the configuration by running the command:

```
refresh -s inetd
```

8.5.4.6 Synchronizing the Cluster

You *must* synchronize the cluster and the node environment before starting HACMP. This copies the cluster definition from one of the remaining cluster nodes to the newly updated node.

Note

You should perform the following steps on one of the remaining configured cluster nodes. Cluster synchronization is a push operation. The configuration of the node where you execute the command will be copied to all other cluster nodes. If you execute the synchronize commands on the newly migrated (and unconfigured) node, you will wipe out the configuration of all nodes in the cluster.

Synchronizing the cluster is broken into two parts:

1. Synchronizing the Cluster Environment

   This step copies the definitions of the cluster topology, including the node and adapter definitions to the updated node.

   ```
   smit hacmp
     Manage Cluster Environment
       Synchronize all Cluster Nodes
   ```

   Examine the configuration on the migrated node by browsing through the SMIT menus. In some instances we had to perform the synchronization step twice to completely configure the cluster environment.

2. Synchronizing the Cluster Nodes

   This copies the configuration of the resources and resource groups to the migrated node.

   ```
   smit hacmp
     Manage Node Environment
       Sync Node Environment
   ```

   Again, you should examine the configuration in the SMIT menus to confirm that it has copied correctly.

8.5.4.7 Restoring the ODM Event Database

If you use customized event scripts in your cluster, and you saved the event ODM database in 8.5.4.2, “Preparations for Upgrading” on page 227, you should now recover the event database using the command:

```
cp /etc/objrepos/HACMPevent.save /etc/objrepos/HACMPevent
```

8.5.4.8 Verifying the Cluster Configuration

When you have finished synchronizing your cluster, you should run the cluster verification procedure to confirm that the cluster is configured correctly. The cluster verification can be run through SMIT, or by running the command `clverify`.

- Verifying the cluster configuration using SMIT

  The path through the SMIT menus depends on the level of HACMP installed, however the SMIT panel is the same in either case:

  – HACMP Version 3.1.1

  ```
  smit hacmp
    Verify Environment
  ```

  – HACMP Version 4.1.1

  ```
  smitty hacmp
    Cluster Configuration
      Cluster Verification
  ```

```
                              Verify Environment

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

                                                    [Entry Fields]
  Verify Cluster Topology, Resources, or Both        both                    +
  Error Count                                       []                        #

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

- Verifying the cluster configuration using the `clverify` utility

  When run directly from the command line, the /usr/sbin/cluster/diag/clverify utility provides a quite unique menu driven user interface to allow the verification of the cluster. It includes two basic options—software and cluster. Each subcommand has further options, which select different verification programs. For example:

```
------------------------------------------------------
 To get help on a specific option, type: help <option>
 To return to previous menu, type: back
 To quit the program, type: quit
------------------------------------------------------
Valid Options are:
topology
config

clverify.cluster>
```

Note

Under HACMP Version 4.1.1, the value returned by clverify more accurately represents the success or failure of the verification. Under Version 3.1.1, the script would return a zero (indicating a success) if the verification script ran to completion, regardless of the actual verification results. In HACMP, the return code from clverify will indicate the results of the actual verification. This is particularly relevant when running the command under SMIT. Under HACMP Version 3.1.1, it was necessary to page through the complete command output to see the results of the verification. Under Version 4.1.1, the result should be more accurately represented by the result (OK or Failed) at the top of the SMIT panel.

8.5.5 Starting HACMP V4.1.1 on the Upgraded Node

The upgrade is complete when you have performed each of the steps identified above. At an appropriate time, HACMP can be restarted on the migrated node using SMIT:

```
smit hacmp
  Cluster Services
    Start Cluster Services
```

```
                              Start Cluster Services

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

                                                      [Entry Fields]
* Start now, on system restart or both                now                      +

  BROADCAST message at startup?                       true                     +
  Startup Cluster Lock Services?                      false                    +
  Startup Cluster Information Daemon?                 false                    +

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

The node will re-acquire its resources and resume providing services to the users. You should now perform whatever testing is possible in your environment (probably not very much).

8.5.6 Upgrading the Next Node

It is now time to decide whether to run in a heterogeneous environment, or to proceed immediately with upgrading the other nodes in your cluster. Before making this decision, you should carefully examine 8.6, “HACMP Version Compatibility.”

To upgrade the other cluster nodes, continue one at a time, following the same procedure as the first node.

8.6 HACMP Version Compatibility

To examine the compatibility and interoperability between HACMP Version 3.1.1 and 4.1.1, we tested the following operations:

- Synchronizing the node configuration
- Verifying the node environment

These steps were tested successfully, and are described in 8.5.4.6, “Synchronizing the Cluster” on page 233 and 8.5.4.8, “Verifying the Cluster Configuration” on page 234. They form an important part of the version compatibility features. The capacity for interoperation between the global ODM of each version allows the node-by-node upgrade of a cluster rather than the all-at-once approach that was required when moving between previous versions.

- Normal operation without faults

It is a basic requirement for heterogeneous operations that nodes at the different levels must be able to interoperate under normal circumstances without faults. This means that the heartbeat or keepalive packets must be successfully

exchanged, and the different levels of cluster manager must agree on the status of the cluster. Again this requirement is met with no problems.

- Adapter swapping
- IP address takeover
- Disk resource takeover
- Customized event scripts

The last four items are described in more detail below.

8.6.1.1 Adapter Swapping

In the event of a LAN adapter failure, if the node has a standby adapter defined on the same network, the IP and hardware addresses of the failing adapter will be swapped onto the standby adapter. This process is known as an adapter swap. An adapter swap will typically take around three seconds for a ethernet adapter, or six seconds for a token-ring adapter.

8.6.1.2 IP Address Takeover

An IP Address Takeover is performed as part of the recovery steps for a complete node failure. In this case, the IP and hardware addresses of the service adapter on the failing node are taken over by a standby adapter on another node in the cluster.

When IP address takeover is first configured on a node, by configuring the service adapter as a resource, the /etc/rc.net and /etc/inittab files are modified by HACMP to give HACMP control over the starting of the addresses that are active on each available adapter. For more information on the modifications to these files, see “TCP/IP Setup” on page 223.

For both adapter swapping and IP address takeover, the following network options should be set to 0 (false) using the no command:

ipforwarding This setting specifies whether the kernel should forward packets. The default value is 0, which specifies to not forward packets. Forwarding is enabled by setting this value to 1.

ipsendredirects Specifies whether the kernel should send IP redirect signals. The default value, 1, specifies that redirects should be sent. This can be disabled by setting the value to 0.

subnetsarelocal Determines if a packet address is on the local network. This option is used in the in_localaddress subroutine. The default value of 1 specifies that addresses that match the local network mask are local. If the value is 0, only addresses matching the local subnetwork are local.

If these values are improperly set, they can cause unpredictable results. The subnetsarelocal value is set automatically by HACMP in the cluster startup script /usr/sbin/cluster/etc/rc.cluster. The other two values can be set in the file /etc/rc.net using the following commands:

```
no -o ipforwarding=0
no -o ipsendredirects=0
```

Note, these two variables both default to zero in any case, however this will not cause any problems and can serve as a reminder of the correct state.

In our testing, adapter swapping functioned correctly between two adapters on a Version 3.1.1 or a Version 4.1.1 system in a heterogeneous cluster. IP address takeover was also successful in a complete node failure scenario, where the IP address is assumed by a standby adapter on a different node in the cluster.

8.6.1.3 Disk Resource Takeover

There is a major problem with disk takeover from a system running AIX Version 4.1 to one running AIX Version 3.2.

Disk takeover involves one or more disks that are physically connected to two or more systems either by a shared SCSI bus, or by multiple serial links. The disks, and the volume groups residing on the disks, are defined on all connected systems. Under normal operations, the volume groups will be varied on to one system only. If that system fails, another node connected to the disks can vary on the volume group or groups, check the file systems to back out any incomplete disk accesses, then begin providing access to the disk resources.

AIX uses a file system journal, called a Journalled File System log (jfslog), to record changes to the structure of the file system. This uses database techniques to log any changes to directories, i-nodes and indirect blocks. When a filesystem is mounted, the fsck command is run to verify the structure of the filesystem. If the filesystem was unmounted cleanly, no action will be taken. If the filesystem is dirty; for example, if the node suffered a power failure, the fsck command will use the journal to determine any filesystem structural changes that were in process, and back out the changes. This ensures that the structure of the filesystem is always consistent.

The problem here arises from the fact that the format of the jfslog changes between AIX Version 3.2 and Version 4.1. This was necessary because of the many enhancements made to filesystem support in the new version, such as filesystem compression and fragments.

The jfslog is upwards compatible. This means that an AIX Version 4.1 system can successfully read and process the jfslog when taking over disks from a Version 3.2 node. However, the reverse is not true. If an AIX Version 3.2 node attempts to take over disk resource from an AIX Version 4.1 node, and the log files are not empty, the file systems will not mount. If you attempt to run the fsck command on the filesystem in this situation, you will see errors such as:

```
log redo processing for /dev/rsharedlv
log redo:not a log file /dev/rsharedlv
failure reply log:=0
```

In this situation, the following manual steps will be required to make the disks available on the new system:

1. Recover from script failure

   If this problem has occurred as the result of an attempted disk takeover in a node failure situation, the cluster manager will still be waiting for the takeover to complete. You will begin seeing additional events, and error messages such as:

```
The cluster has been in reconfiguration too long
```

You should stop these errors by recovering from the script error. Use the SMIT panel:

```
smit hacmp
  Cluster Recovery Aids
    Recover From Script Failure
```

Select the service adapter for the node that was attempting to takeover the disk resources and press Enter.

2. Reinitialize the log

   It is necessary to reinitialize the log file. The name of the jfslog associated with the file systems in a volume group can be shown with the `lsvg -l <vgname>` command. For example:

```
# lsvg -l sharedvg
sharedvg:
LV NAME             TYPE       LPs   PPs  PVs  LV STATE      MOUNT POINT
sharedloglv         jfslog     1     2    2    closed/syncd  N/A
sharedlv            jfs        5     10   2    closed/syncd  /sharedfs
```

   The log file (or files - there could be more than one per volume group) are initialized using the `logform` command:

```
# logform /dev/sharedloglv
logform: destroy /dev/sharedloglv (y)?y
```

   This will clear any existing entries in the log.

3. Check the file systems

   The file systems will still be marked as dirty. Thus it is necessary to check them with the `fsck` command. The `fsck` command will attempt to correct any problems it finds in the filesystem; however, without the log file, it is unable to fix some of the possible problems.

   Attention!

   There is a significant possibility that data will be lost during this operation.

   The `fsck` command is run as follows:

```
# fsck /dev/sharedlv

** Checking /dev/sharedlv (/share)
** Phase 0  - Check Log
log redo processing for /dev/sharedlv
** Phase 1  - Check Blocks and Sizes
** Phase 2  - Check Pathnames
** Phase 3  - Check Connectivity
** Phase 4  - Check Reference Counts
** Phase 5  - Check Inode Map
** Phase 6  - Check Block Map
The superblock is marked dirty.; FIX? y
289 files 3840 blocks 37120 free
***** File system was modified *****
```

4. Mount the file systems

For this reason, we recommend that the version compatibility functions be used only as an aid to migration, and not as a long-term solution.

8.6.1.4 Customized Event Scripts

There are two factors relating to customized event scripts that should be considered when migrating to HACMP Version 4.1.1.

The first consideration is that the events are lost during the process, and must be saved by the user and restored after the upgrade. This is described in 8.5.4.2, “Preparations for Upgrading” on page 227 and 8.5.4.7, “Restoring the ODM Event Database” on page 233.

The second consideration is that the handling and definition of events changes slightly between versions. In HACMP Version 3.1.1, a customized event script is defined for a specific event, on a specific node in the cluster. Separate event databases (HACMPevent) are kept on each node. When a new customized event is added, it is added only to the database on the selected node using the global ODM.

In HACMP Version 4.1.1, a customized script is defined simply for a specific event in the cluster. Thus the event script (if it exists and is executable) will be executed on all nodes in the cluster. The event ODM databases are synchronized between the clusters.

For this reason, if you wish to have different customized event scripts run on specific nodes in a cluster, it is necessary to add additional coding to the event script to determine the node and perform the appropriate actions. This additional coding can utilize the environment variable LOCALNODENAME which is set by HACMP and exported to all event scripts.

An example of the required logic is shown below:

```
case $LOCALNODENAME in
mickey) echo "Running on node mickey" >>/tmp/myscript.log
        # Do processing for node mickey
        ;;
goofy)  echo "Running on node goofy" >>/tmp/myscript.log
        # Do processing for node goofy
        ;;
*)      echo "Error: Unknown node $LOCALNODENAME" >>/tmp/myscript.log
        ;;
esac
```

Chapter 9. Sample HACMP Migration

This chapter illustrates the HACMP migration process by following a sample migration that we performed as part of our testing.

9.1 Sample Environment

As an example, we upgraded a cluster consisting of two RISC System/6000 Model 530s running AIX V3.2.5 and HACMP V3.1.1 to AIX V4.1.4 and HACMP V4.1.1. Below is the architecture of the cluster in the initial testing environment. The two nodes are named mickey and goofy.

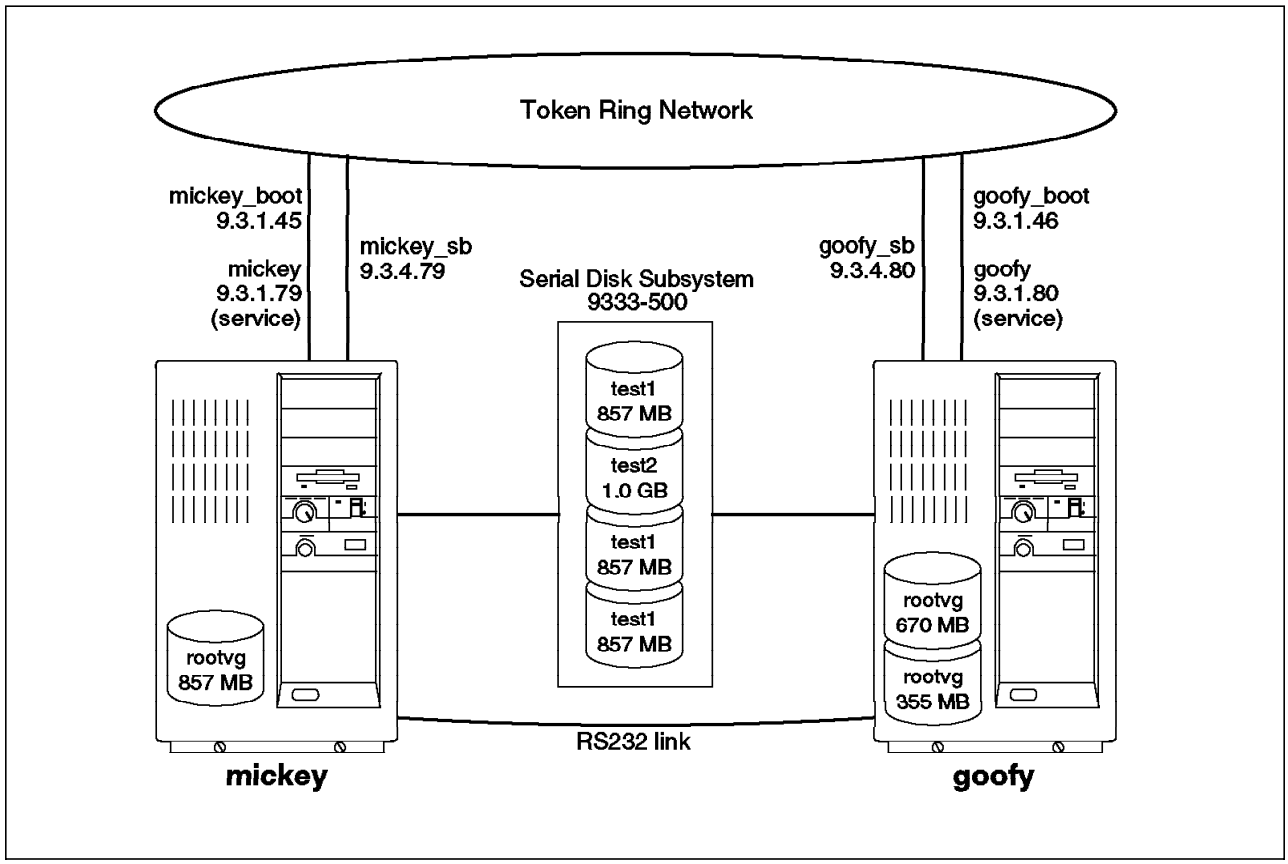

Figure 15. Initial HACMP configuration

9.2 Initial Configuration

Before starting our migration, we documented our cluster as follows:

9.2.1 HACMP Level

We verified the level of HACMP installed on the system using the command `lslpp -L "cluster*"`. Both systems have the same software installed.

```
Processing.....Please Wait.
 Description                                             State     Fix Id
 ------------------------------------------------------  ------    ------------

 cluster.client 3.1.0.0
   No Maintenance Level Applied.
     HACMP/6000                                            A     U439890

 cluster.clvm 3.1.0.0
   No Maintenance Level Applied.

 cluster.server 3.1.0.0
   No Maintenance Level Applied.
     HACMP/6000                                            A     U439890


State Codes:
 A -- Applied.
 B -- Broken.
 C -- Committed.
 N -- Not Installed, but was previously installed/seen on some media.
 - -- Superseded, not Applied.
 ? -- Inconsistent State...Run lppchk -v.
```

We can see that we have HACMP Version 3 Release 1. In addition, since there is a PTF installed, we know that we are on Modification level 1. Thus we are using HACMP Version 3.1.1.

Note, although the Concurrent LVM (CLVM) is installed on these systems, it was not used in our testing. We did not test any concurrent access operations.

9.2.2 Hostname

Using the command `hostname`, we verified the names of our two nodes.

On mickey:

```
mickey
```

On goofy:

```
goofy
```

9.2.3 Cluster ID and Name

The cluster name and ID were obtained from the command `cllsclstr`.

```
ID      Name

1       disney
```

9.2.4 Node Names

In our case, the cluster node names used are the same as the hostnames. The command clnodename shows the node names as follows.

```
goofy
mickey
```

9.2.5 Network Adapters

From the SMIT menus we can see the adapters defined to our cluster are shown on the screen below.

```
                                Configure Adapters

Move cursor to desired item and press Enter.

  Add an Adapter
  Change / Show an Adapter
  Remove an Adapter
  +-------------------------------------------------------------------------+
  |                            Adapter to Change                            |
  |                                                                         |
  | Move cursor to desired item and press Enter.                            |
  |                                                                         |
  |   goofy                                                                 |
  |   goofy_boot                                                            |
  |   goofy_sb                                                              |
  |   goofy_tty0                                                            |
  |   mickey                                                                |
  |   mickey_boot                                                           |
  |   mickey_sb                                                             |
  |   mickey_tty0                                                           |
  |                                                                         |
  | F1=Help                F2=Refresh             F3=Cancel                 |
  | F8=Image               F10=Exit               Enter=Do                  |
F1| /=Find                 n=Find Next                                      |
F9+-------------------------------------------------------------------------+
```

The detailed definition of each adapter is shown below.

Service Adapter: mickey

```
 Adapter IP Label                                 mickey
 New Adapter Label                               []
 Network Type                                    [token]
 Network Name                                    [trnet1]
 Work Attribute                                   public
 Adapter Function                                 service
 Adapter Identifier                              [9.3.1.79]
 Adapter Hardware Address                        [0x42005aa8b484]
 Node Name                                       [mickey]
```

Boot Adapter: mickey_boot

```
Adapter IP Label                                 mickey_boot
New Adapter Label                               []
Network Type                                    [token]
Network Name                                    [trnet1]
Work Attribute                                   public
Adapter Function                                 boot
Adapter Identifier                              [9.3.1.45]
Adapter Hardware Address                        []
Node Name                                       [mickey]
```

Standby Adapter: mickey_sb

```
Adapter IP Label                                 mickey_sb
New Adapter Label                               []
Network Type                                    [token]
Network Name                                    [trnet1]
Work Attribute                                   public
Adapter Function                                 standby
Adapter Identifier                              [9.3.4.79]
Adapter Hardware Address                        []
Node Name                                       [mickey]
```

Serial Adapter: mickey_tty0

```
Adapter IP Label                                 mickey_tty0
New Adapter IP label                            []
Network Type                                    [rs232]
Network Attribute                                serial
Adapter Function                                 service
Adapter Identifier                              [/dev/tty0]
Adapter Hardware Address                        []
Node Name                                       [mickey]
```

Service Adapter: goofy

```
Adapter IP Label                                 goofy
New Adapter IP Label                            []
Network Type                                    [token]
Network Name                                    [trnet1]
Network Attribute                                public
Adapter Function                                 service
Adapter Identifier                              [9.3.1.80]
Adapter Hardware Address                        [0x42005aa8d1f3]
Node Name                                       [goofy]
```

Boot Adapter: goofy_boot

```
Adapter IP Label                                 goofy_boot
New Adapter IP Label                            []
Network Type                                    [token]
Network Name                                    [trnet1]
Network Attribute                                public
Adapter Function                                 boot
Adapter Identifier                              [9.3.1.46]
Adapter Hardware Address                        []
Node Name                                       [goofy]
```

Standby Adapter: goofy_sb

```
Adapter IP Label                                 goofy_sb
New Adapter IP Label                            []
Network Type                                    [token]
Network Name                                    [trnet1]
Network Attribute                                public
Adapter Function                                 standby
Adapter Identifier                              [9.3.4.80]
Adapter Hardware Address                        []
Node Name                                       [goofy]
```

Serial Adapter: goofy_tty0

```
Adapter IP Label                                 goofy_tty0
New Adapter IP Label                            []
Network Type                                    [rs232]
Network Name                                    [serial1]
Network Attribute                                serial
Adapter Function                                 service
Adapter Identifier                              [/dev/tty0]
Adapter Hardware Address                        []
Node Name                                       [goofy]
```

9.2.6 Network Interfaces

We used the commands `ifconfig tr0` and `ifconfig tr1` to examine the network configuration of our systems. Note, you will get different results running this command with or without HACMP running. The following commands were run while HACMP was running, and thus the service address appears on the first adapter. When HACMP is not running, you will instead see the boot address.

On mickey:

```
tr0: flags=8063<UP,BROADCAST,NOTRAILERS,RUNNING,ALLCAST>
        inet 9.3.1.79 netmask 0xffffff00 broadcast 9.3.1.255

tr1: flags=8063<UP,BROADCAST,NOTRAILERS,RUNNING,ALLCAST>
        inet 9.3.4.79 netmask 0xffffff00 broadcast 9.3.4.255
```

On goofy:

```
tr0: flags=8063<UP,BROADCAST,NOTRAILERS,RUNNING,ALLCAST>
        inet 9.3.1.80 netmask 0xffffff00 broadcast 9.3.1.255

tr1: flags=8063<UP,BROADCAST,NOTRAILERS,RUNNING,ALLCAST>
        inet 9.3.4.80 netmask 0xffffff00 broadcast 9.3.4.255
```

Note, in each case, the two adapters are on different TCP/IP subnets.

9.2.7 Serial Interfaces

We can see the configuration of the serial ports used to provide a non-IP connection between the systems by using the command `lsattr -El tty0`.

```
speed           9600       BAUD rate
kill            ^u         KILL character
erase           ^h         ERASE character
eof             ^d         END OF FILE character
dsusp           ^y         DELAY SUSPEND PROCESS character
susp            ^z         SUSPEND PROCESS character
lnext           ^v         LITERAL NEXT character
start           ^q         START character
stop            ^s         STOP character
werase          ^w         WORD ERASE character
parity          none       PARITY
bpc             8          BITS per character
stops           1          Number of STOP BITS
xon             yes        XON-XOFF handshaking
term            dumb       TERMINAL type
login           disable    Enable LOGIN
runmodes        hupcl,cread,brkint,icrnl,opost,tab3,onlcr,isig,icanon,
                echo,echoe,echok,echoctl,echoke,imaxbel,iexten
                           STTY attributes for RUN TIME
quit            ^\         QUIT character
intr            ^c         INTERRUPT character
reprint         ^r         REPRINT LINE character
logmodes        hupcl,cread,echoe,cs8,ixon,ixoff
                           STTY attributes for LOGIN
eol             ^@         END OF LINE character
eol2            ^          2nd END OF LINE character
discard         ^o         DISCARD character
autoconfig      available  STATE to be configured at boot time
imap            none       INPUT map file
omap            none       OUTPUT map file
csmap           sbcs       CODESET map file
tbc             64         TRANSMIT buffer count
timeout         0          TIME before advancing to next port setting
shell           no         RUN shell activity manager
logger                     Optional LOGGER name
2200flow        disable    2200 Flow Control
2200print       disable    2200 Print Control
altpin          disable    Use Alternate RJ-45 Pinouts
bufsize         100        Transparent Print Printer Buffer Size
edelay          100        Receive Event Delay Time
fastcook        enable     Perform Cooked Processing in Adapter
forcedcd        disable    Force Carrier
maxchar         50         Transparent Print Maximum Character Packet Size
maxcps          100        Transparent Print Maximum Characters per Second
offstr          \033[4i    Transparent Print OFF String
onstr           \033[5i    Transparent Print ON String
rtrig           3          Read Trigger
rts             no         RTS-CTS handshaking
ttyprog_action  respawn    N/A
ttyprog_rlevel  2          Run level
xprint_off_str  \033[4i    Printer OFF string
xprint_on_str   \033[5i    Printer ON string
xprint_priority 30         Priority of transparent print
```

This configuration is common between both systems, and has been reformatted to fit this redbook.

9.2.8 Shared Disk

We then examined the shared disk configuration on the systems using the command `lsdev -Cc disk`

On mickey:

```
hdisk0 Available 00-08-00-00 1.0 GB SCSI Disk Drive
hdisk1 Available 00-03-00-00 857MB Serial-Link Disk Drive
hdisk2 Available 00-03-00-01 1.07GB Serial-Link Disk Drive
hdisk3 Available 00-03-00-02 857MB Serial-Link Disk Drive
hdisk4 Available 00-03-00-03 857MB Serial-Link Disk Drive
```

This output shows us that there is one 857 MB disk connected to the SCSI adapter in slot 8. This disk is internal within the 530 cabinet and is used for the root volume group in this configuration. The four serial link disks are connected to a serial disk adapter in slot 3.

On goofy:

```
hdisk0 Available 00-08-00-00 670 MB SCSI Disk Drive
hdisk1 Available 00-08-00-20 320 MB SCSI Disk Drive
hdisk2 Available 00-06-00-00 857MB Serial-Link Disk Drive
hdisk3 Available 00-06-00-01 1.07GB Serial-Link Disk Drive
hdisk4 Available 00-06-00-02 857MB Serial-Link Disk Drive
hdisk5 Available 00-06-00-03 857MB Serial-Link Disk Drive
```

Goofy has two internal SCSI disks connected to the SCSI adapter in slot 8. We can see that the serial disk adapter in goofy is in slot 6, but is connected to the same four disks as mickey.

Note, in our environment it was not necessary to examine the addresses of the SCSI adapters as we did not have any shared SCSI resources.

9.2.9 Volume Groups

Using the command `lspv`, we displayed the volume groups and their distribution across the available disks.

On mickey:

```
hdisk0          000111874109e674    rootvg
hdisk1          0000411925a74610    test1
hdisk2          0000411979d1f8af    test2
hdisk3          00002819699e632f    test1
hdisk4          000005080b85c688    test1
```

On goofy:

```
hdisk0          000122180006c324    rootvg
hdisk1          0002479088f5f347    rootvg
hdisk2          0000411925a74610    test1
hdisk3          0000411979d1f8af    test2
hdisk4          00002819699e632f    test1
hdisk5          000005080b85c688    test1
```

Volume groups test1 and test2 are on the shared 9333 disks.

We can use the commands `lsvg test1` on mickey, and `lsvg test2` on goofy to examine the configuration of these volume groups.

On mickey:

```
VOLUME GROUP:   test1                   VG IDENTIFIER:  00014732b5a91022
VG STATE:       active                  PP SIZE:        4 megabyte(s)
VG PERMISSION:  read/write              TOTAL PPs:      609 (2436 megabytes)
MAX LVs:        256                     FREE PPs:       548 (2192 megabytes)
LVs:            2                       USED PPs:       61 (244 megabytes)
OPEN LVs:       2                       QUORUM:         1
TOTAL PVs:      3                       VG DESCRIPTORS: 3
STALE PVs:      0                       STALE PPs       0
ACTIVE PVs:     3                       AUTO ON:        no
```

On goofy:

```
VOLUME GROUP:   test2                   VG IDENTIFIER:  00014732ca66234e
VG STATE:       active                  PP SIZE:        4 megabyte(s)
VG PERMISSION:  read/write              TOTAL PPs:      255 (1020 megabytes)
MAX LVs:        256                     FREE PPs:       228 (912 megabytes)
LVs:            3                       USED PPs:       27 (108 megabytes)
OPEN LVs:       2                       QUORUM:         1
TOTAL PVs:      1                       VG DESCRIPTORS: 2
STALE PVs:      0                       STALE PPs       0
ACTIVE PVs:     1                       AUTO ON:        no
```

Note, the AUTO ON field is set to no. In HACMP it is important that the shared volume groups are not varied on automatically by the operating system. This allows HACMP to control access to the volume groups, varying them on to the owner under normal circumstances, or to backup nodes when appropriate.

9.2.10 File Systems on the Shared Disk

Using the command `lsvg -l test1` and `lsvg -l test2` we displayed the information about file systems in the non-rootvg volume groups.

On mickey:

```
test1:
LVNAME              TYPE       LPs   PPs   PVs  LV STATE      MOUNT POINT
loglvtest1          jfslog     1     1     1    open/syncd    N/A
lvtest1             jfs        20    60    3    open/syncd    /test1
```

Note that the lvtest1 Logical Volume (LV) has 20 Logical Partitions (LPs) but 60 Physical Partitions (PPs) on 3 Physical Volumes (PVs). This shows that 3 mirrored copies of this LV (and thus the /test1 filesystem) are being maintained.

On goofy:

```
test2:
LVNAME              TYPE       LPs   PPs   PVs  LV STATE      MOUNT POINT
loglvtest2          jfslog     1     1     1    open/syncd    N/A
lvtest2             jfs        25    25    1    open/syncd    /test2
```

In addition, we also checked the device major numbers used for each volume group using the command `ls -l /dev/test*`. The device major numbers for each volume group are the same on both systems.

```
crw-rw----   1 root     system    26,  0 May 07 19:30 /dev/test1
crw-rw----   1 root     system    27,  0 May 07 19:28 /dev/test2
```

That is, volume group test1 uses a device major number of 26 on both nodes. Volume group test2 uses a device major number of 27.

9.2.11 Resource Groups

We have two resource groups defined in our cluster. The resource groups are named mickeyrg and goofyrg. This is intended to indicate that in normal operations, the resources in the resource group mickeyrg will be accessed by node mickey. Similarly, the resources in group goofyrg will be accessed by node goofy under normal operations.

The first resource group:

```
 Resource Group Name                                 mickeyrg
 Node Relationship                                   cascading
 Participating Node Names                            mickey goofy

 Service IP label                                   [mickey]
 Filesystems                                        [/test1]
 Filesystems to Export                              [/test1]
 Filesystems to NFS mount                           []
 Volume Groups                                      []
 Concurrent Volume groups                           []
 Raw Disk PVIDs                                     []
 Application Servers                                []
 Miscellaneous Data                                 []

 Inactive Takeover Activated                         false
 9333 Disk Fencing Activated                         false
 SSA Disk Fencing Activated                          false
```

The second resource group:

```
  Resource Group Name                                  goofyrg
  Node Relationship                                    cascading
  Participating Node Names                             goofy mickey

  Service IP label                                    [goofy]
  Filesystems                                         [/test2]
  Filesystems to Export                               [/test2]
  Filesystems to NFS mount                            []
  Volume Groups                                       []
  Concurrent Volume groups                            []
  Raw Disk PVIDs                                      []
  Application Servers                                 []
  Miscellaneous Data                                  []

  Inactive Takeover Activated                          false
  9333 Disk Fencing Activated                          false
  SSA Disk Fencing Activated                           false
```

9.2.12 Application Servers

We are not using any application servers in our sample cluster.

9.2.13 User Defined Cluster Events

Our cluster includes three customized event scripts for the event node_down_remote. This customization is performed on the node mickey only. That is, when the node mickey detects the remote node (goofy) has failed, it will execute the three customized event scripts shown below. In this very basic example, the scripts simply echo a timestamp and the parameters passed to the script into a file.

```
                          Change/Show Cluster Events

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

                                                      [Entry Fields]
  Node Name                                          mickey

  Event Name                                         node_down_remote

  Description                                        Script run when it is >

  Event Command                                     [/usr/sbin/cluster/even>

  Notify Command                                     [/usr/local/note-event]
  Pre-event Command                                  [/usr/local/pre-event]
  Post-event Command                                 [/usr/local/post-event]
  Recovery Command                                   []
  Recovery Counter                                   [0]                    #


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

The three event scripts are shown below:

- note-event

```
#!/bin/ksh

date >> /usr/local/note-event.args
echo $* >> /usr/local/note-event.args
```

- pre-event

```
#!/bin/ksh

date >> /usr/local/pre-event.args
echo $* >> /usr/local/pre-event.args
```

- post-event

```
#!/bin/ksh

date >> /usr/local/post-event.args
echo $* >> /usr/local/post-event.args
```

9.2.14 Run Time Parameters

The run time parameters are shown in the SMIT panel below.

```
                        Configure Run Time Parameters

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

                                                      [Entry Fields]
  Node Name                                          mickey

  Debug Level                                        low                  +
  Host uses NIS or Name Server                       true                 +









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

9.2.15 Clients

In our sample migration, we are considering only the cluster nodes and are not discussing client systems.

9.3 Testing Procedures

Before starting the migration, we first verified that the cluster was operating correctly. This testing included simulating:

- Network adapter failure
- Network failure

- Full node failure

These tests were then later repeated to ensure correct interoperation between Versions 3.1.1 and 4.1.1, and the correct operation of the final Version 4.1.1 cluster. The testing procedures used are discussed in detail in 8.4, “Determining Test Procedures” on page 212.

9.4 Migration Process

In our environment all migration testing was performed in HACMP Mode 2. We did not test any concurrent access operations. HACMP V3.1.1 was run on AIX V3.2.5 and this was upgraded to HACMP V4.1.1 running on AIX V4.1.4.

We chose to first perform the upgrade on node goofy. We then tested to verify that HACMP V4.1.1 on goofy can work together with HACMP V3.1.1 on mickey. Finally we upgraded node mickey.

The upgrade stages are shown below:

1. Stopping HACMP on goofy

   By stopping HACMP on goofy with the takeover option, we allowed mickey to take over the resources owned by goofy.

2. Backing up goofy
3. Upgrading the operating system on goofy

   After upgrading AIX, we took the required manual steps to access the shared volume groups, set the volume group device major numbers, and to allow us to reboot the system with TCP/IP running.

4. Upgrading HACMP on goofy

   After upgrading HACMP to Version 4.1.1, we then reconfigured the node by synchronizing from mickey, and recovered the customized event configuration.

5. Starting HACMP V4.1.1 on goofy

   We then started HACMP on goofy, allowing it to take back its resources from mickey.

6. Testing of HACMP V3.1.1 and HACMP V4.1.1 interoperation
7. Stopping HACMP on mickey

   This time, goofy was allowed to take over the resources of mickey.

8. Backing up mickey
9. Upgrading the operating system on mickey
10. Upgrading HACMP on mickey
11. Starting HACMP V4.1.1 on mickey

    By starting HACMP V4.1.1 on node mickey, we allowed mickey to regain its resources.

12. Verify the correct operation of the Version 4.1.1 cluster.

9.4.1 Stopping HACMP on goofy

We stopped HACMP from SMIT, using the graceful with takeover option, and watched the hacmp.out and cluster.log files. Once all event processing was complete we verified that the disk and IP address takeovers were successful by running the `df` and `netstat -i` commands on mickey.

```
# df
Filesystem    Total KB    free %used   iused %iused Mounted on
/dev/hd4         12288    3824   68%     947    23% /
/dev/hd9var       4096    2540   37%     112    10% /var
/dev/hd2        552960   55612   89%   20233    14% /usr
/dev/hd3         20480   16340   20%      49     0% /tmp
/dev/hd1          4096    3932    4%      17     1% /home
/dev/lvtest1     81920   79312    4%      16     1% /test1
/dev/lvtest2    102400   99024    4%      16     1% /test2

# netstat -i
Name  Mtu   Network     Address         Ipkts    Ierrs Opkts    Oerrs Coll
lo0   1536  <Link>                         2903     0     2903     0     0
lo0   1536  127         localhost          2903     0     2903     0     0
tr0   1492  <Link>                         2215     0     2020     0     0
tr0   1492  9.3.1       mickey.itsc.aus    2215     0     2020     0     0
tr1   1492  <Link>                         1081     0       80     0     0
tr1   1492  9.3.1       goofy.itsc.aust    1081     0       80     0     0
```

We can see that both the filesystem normally owned by node mickey (/test1) and the filesystem normally owned by node goofy (/test2) are available on node mickey, and that adapter tr1 has taken over goofy's service address.

9.4.2 Backing Up goofy

After stopping HACMP, we made a bootable system backup using the `mksysb` command.

However, before stopping HACMP, we also saved the /test1 filesystem. The /test1 filesystem is the only non-rootvg filesystem that is accessed by goofy under fault free operations. This filesystem was backed up using the command:

```
cd /
tar -cvf /dev/rmt0 ./test1
```

This backup of /test1 was made before stopping HACMP so that the filesystem would still be available on goofy before being taken over by mickey.

9.4.3 Upgrading AIX On goofy

For details on upgrading the AIX operating system, see *A Holistic Approach to AIX V4.1 Migration, Volume 1*, SG24-4652. We chose to perform a migration installation as this method requires the least reconfiguration after the installation.

After upgrading AIX, we considered the various problems discussed in 8.5.3.1, "Problems Encountered After the AIX Upgrade" on page 215.

9.4.3.1 Volume Groups on Shared Disks

Because the volume groups test1 and test2 on the shared disks were being accessed by mickey during the migration, they could not be imported by goofy at the end of the migration process. Thus the volume groups were not defined on goofy after the migration.

9.4.3.2 Device Major Numbers

Before the migration, the volume groups test1 and test2 used device major numbers of 26 and 27 respectively. After the migration, these values were no longer available. We checked the available values on each node using the `lvlstmajor` command. In each case, the output was as follows:

```
35...
```

This indicates that the values 35 and above are available to use as device major numbers; however, to avoid the likelihood of having to change the value again in the future, we selected values of 60 and 61 for our volume groups. We decided that since we would be performing testing on the heterogeneous cluster, we should change the major number value on both mickey and goofy to keep the value the same on both systems.

We used the following procedure to both recover the configuration of the shared volume groups, and, at the same time, to change the value used for the device major numbers on each system.

1. We kicked our imaginary clients and users off the system :-)

   While frivolous, this step is here to make the point that the filesystems and applications will not be available to your users during this procedure.

2. Unmounted the file systems on mickey

   ```
   # umount /test1
   # umount /test2
   ```

3. Varied off the shared volume groups on mickey

   ```
   # varyoffvg test1
   # varyoffvg test2
   ```

4. Imported the volume groups on goofy

   We imported the volume groups using `smit importvg`:

   ```
                    Import a Volume Group
   TYPE or select values in entry fields.
   Press Enter AFTER making all desired changes.

                                                  [Entry Fields]
   *VOLUME GROUP name                              [test1]
   *PHYSICAL VOLUME name                           [hdisk2]
   ACTIVATE volume group after it is               yes
   imported?
   Volume Group MAJOR NUMBER                      [60]


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

   Similarly for volume group test2:

```
                  Import a Volume Group
TYPE or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                      [Entry Fields]
*VOLUME GROUP name                                     [test2]
*PHYSICAL VOLUME name                                  [hdisk3]
ACTIVATE volume group after it is                      yes
imported?
Volume Group MAJOR NUMBER                             [61]


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

5. Changed the volume group settings to match original volumes

   It is important not to forget to set the Activate volume group AUTOMATICALLY and QUORUM settings for each volume group. In *every* HACMP case, the volume group must be set to *not* activate automatically at boot time. The QUORUM setting may also need to be changed from its default setting of yes.

   The quorum settings are changed through the smit chvg menu panel. We set both volume groups as follows:

```
                            Change a Volume Group

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

                                                        [Entry Fields]
* VOLUME GROUP name                                    test1
* Activate volume group AUTOMATICALLY                  no                      +
    at system restart?
* A QUORUM of disks required to keep the volume        no                      +
   group on-line ?









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

   The same settings were used for the test2 volume group.

6. Varied off the volume groups on goofy

   Do not export the volume groups at this stage.

7. Exported the volume groups on mickey

```
# exportvg test1
# exportvg test2
```

8. Imported the volume groups on mickey

   This step was identical to step 4 on page 254, apart from the fact that the disk names are different. On mickey, we imported volume group test1 from disk hdisk1, and volume group test2 from disk hdisk2. This is a result of the fact that goofy has one additional disk in the root volume group. We again specified the new device major numbers of 60 and 61.

9. Changed the volume group settings on mickey

   This step was again identical to the operation performed on goofy in step 5 on page 255 above.

At the end of this procedure, both shared volume groups were again defined on each system, with the volume group settings correct, and with consistent device major numbers.

Note: In our sample environment we did not need to set the SCSI address or enable target mode SCSI as we did not use any shared SCSI disks.

9.4.3.3 TCP/IP Setup

To avoid problems with the TCP/IP configuration, we deleted the lines from the `/etc/rc.net` script file as described in “TCP/IP Setup” on page 223. After taking this step, we had no problems with TCP/IP.

9.4.4 Upgrading HACMP on Node goofy

Before installing the new version of HACMP software, we first met the prerequisites by installing the fileset bos.compat.lan, and rebooting the system as instructed.

We also saved the event ODM file by copying it:

```
cp /etc/objrepos/HACMPevent /etc/objrepos/HACMPevent.save
```

We then installed the following HACMP Version 4.1.1 software:

```
  Fileset                      Level  State  Description
  ----------------------------------------------------------------------------
  cluster.adt.client.demos   4.1.1.0    C    HACMP Client Demos
  cluster.adt.client.include
                             4.1.1.0    C    HACMP Client Include Files
  cluster.adt.client.samples.clinfo
                             4.1.1.0    C    HACMP Client CLINFO Samples
  cluster.adt.client.samples.clstat
                             4.1.1.0    C    HACMP Client Clstat Samples
  cluster.adt.client.samples.demos
                             4.1.1.0    C    HACMP Client Demos Samples
  cluster.adt.client.samples.libcl
                             4.1.1.0    C    HACMP Client LIBCL Samples
  cluster.adt.server.demos   4.1.1.0    C    HACMP Server Demos
  cluster.adt.server.samples.demos
                             4.1.1.0    C    HACMP Server Sample Demos
  cluster.adt.server.samples.images
                             4.1.1.0    C    HACMP Server Sample Images
  cluster.base.client.lib    4.1.1.0    C    HACMP Base Client Libraries
  cluster.base.client.rte    4.1.1.0    C    HACMP Base Client Runtime
  cluster.base.client.utils  4.1.1.0    C    HACMP Base Client Utilities
  cluster.base.server.diag   4.1.1.0    C    HACMP Base Server Diags
  cluster.base.server.events
                             4.1.1.0    C    HACMP Base Server Events
  cluster.base.server.rte    4.1.1.0    C    HACMP Base Server Runtime
  cluster.base.server.utils  4.1.1.0    C    HACMP Base Server Utilities
  cluster.clvm.rte           4.1.1.0    C    HACMP for AIX Concurrent Access
  cluster.man.en_US.client.data
                             4.1.1.0    C    HACMP Client Man Pages - U.S.
                                             English
  cluster.man.en_US.server.data
                             4.1.1.0    C    HACMP Server Man Pages - U.S.
                                             English
  cluster.msg.en_US.client   4.1.1.0    C    HACMP Client Messages - U.S.
                                             English
  cluster.msg.en_US.server   4.1.1.0    C    HACMP Server Messages - U.S.
                                             English
  cluster.vsm.server         4.1.1.0    C    HACMP Visual System Management
                                             Configuration Utility


State Codes:
 A -- Applied.
 B -- Broken.
 C -- Committed.
 O -- Obsolete.  (partially migrated to newer version)
 ? -- Inconsistent State...Run lppchk -v.
```

Once the new version of software was installed, we edited the file /etc/inetd.conf to remove the duplicate GODM entry as described in 8.5.4.5, “Problems Encountered After the HACMP Upgrade” on page 232. We then synchronized the cluster from the node mickey. This configured goofy for us.

We investigated the customized event configuration, and discovered that as expected, the customized event scripts were not configured.

```
                          Change/Show Cluster Events

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

                                                      [Entry Fields]
  Node Name                                          mickey

  Event Name                                         node_down_remote

  Description                                        Script run when it is >

  Event Command                                     [/usr/sbin/cluster/even>

  Notify Command                                     []
  Pre-event Command                                  []
  Post-event Command                                 []
  Recovery Command                                   []
  Recovery Counter                                   [0]                    #



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

We copied the saved ODM database back to its original name:

```
cp /etc/objrepos/HACMPevent.save /etc/objrepos/HACMPevent
```

Our events were now configured again as before the upgrade.

9.4.5 Verifying the Cluster

We then ran cluster verification from each cluster node. When run from mickey, the verification ran cleanly and succesfully. When run from goofy, the cluster verification complained that the ODM event databases were not the same on the two systems.

This is a consequence of the differences in event handling between the two HACMP versions. For more information, see 8.6.1.4, “Customized Event Scripts” on page 240.

9.4.6 Starting HACMP on goofy

Having completed the configuration of the newly upgraded node, we started HACMP. After waiting for all event processing to complete, we checked that goofy had regained its resources.

```
# df
Filesystem    Total KB    free %used   iused %iused Mounted on
/dev/hd4         12288    4884   60%     945    23% /
/dev/hd9var       4096    2584   36%     112    10% /var
/dev/hd2        552960   56980   89%   20248    14% /usr
/dev/hd3         20480   16568   19%      51     0% /tmp
/dev/hd1          4096    3932    4%      17     1% /home
/dev/lvtest2    102400   99024    4%      16     1% /test2

# netstat -i
Name  Mtu   Network     Address          Ipkts    Ierrs Opkts    Oerrs Coll
lo0   1536  <Link>                          15780     0    15780     0     0
lo0   1536  127         localhost           15780     0    15780     0     0
tr1   1492  <Link>                          73759     0    54954     0     0
tr1   1492  9.3.4       goofy_sb.itsc.a     73759     0    54954     0     0
tr0   1492  <Link>                          77462     0    60015     0     0
tr0   1492  9.3.1       goofy.itsc.aust     77462     0    60015     0     0
```

This shows that the test2 filesystem is now mounted on goofy, and that goofy is once again using both its service and standby adapters on the two token-ring adapters.

9.5 Interoperability Between HACMP V3.1.1 and V4.1.1

Now with goofy running HACMP V4.1.1 on AIX V4.1.4. and mickey running HACMP V3.1.1 on AIX V3.2.5, we were ready to test the interoperability of the two versions in a heterogeneous cluster. The results of our testing are shown below:

9.5.1 Simulated Adapter Failure

The first and simplest tests examined the ability to cope with the failure of a network adapter on one of the nodes. In this situation, adapter swapping is emplyed to take over the IP and hardware addresses of the failing adapter to the standby adapter on the same node. This does not test a takeover between different versions of the operating system, but does verify that the cluster managers agree on the problem that has occurred and will behave accordingly.

To test this function, we simply unplugged the adapter currently being used for the service address of one node. HACMP recognized that the service adapter was no longer available, and changed the service address onto the other adapter in place of the standby address. We then repeated this test to move the service address back to the original adapter, and performed the same tests on the other node.

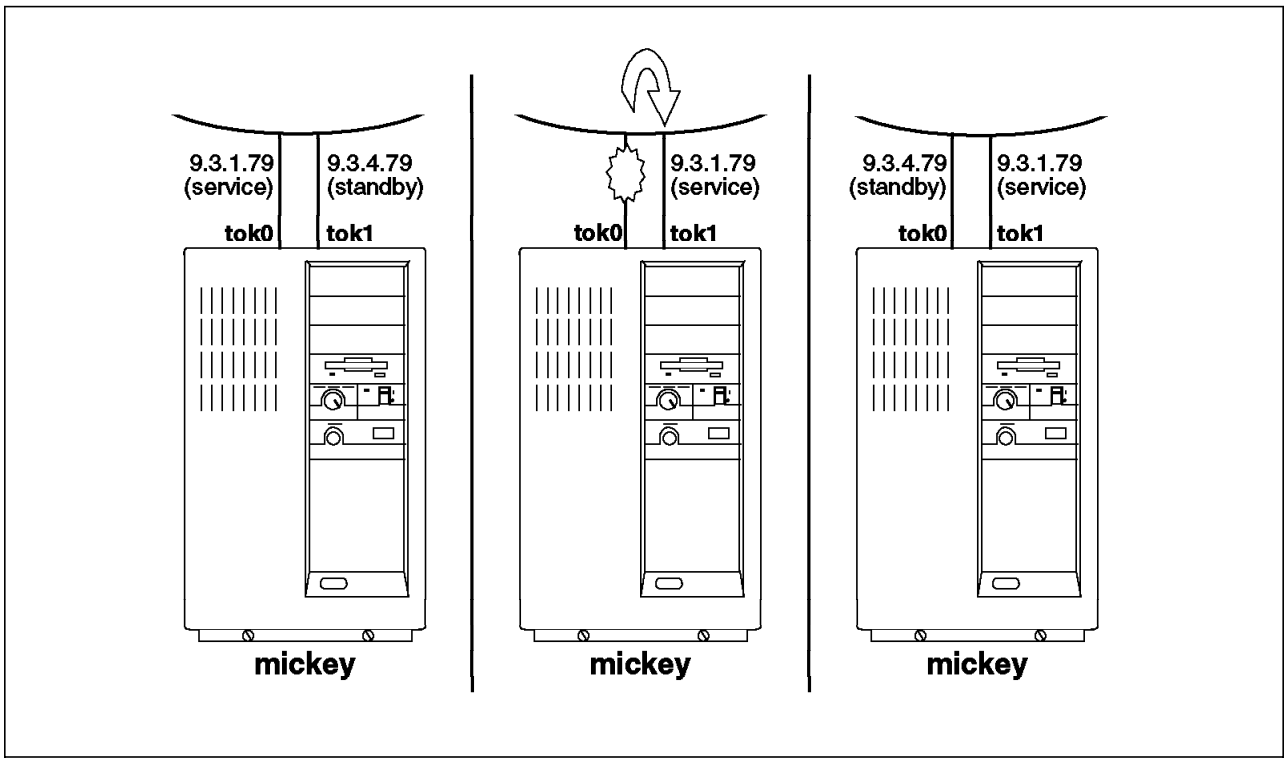

Figure 16. Simulated Service Adapter Failure

Figure 16 shows the sequence of events as the service adapter of mickey is unplugged from the network to simulate an adapter failure. This can be best seen in the output of the `netstat -i` command run on mickey. The output before unplugging is:

```
Name  Mtu   Network     Address         Ipkts    Ierrs Opkts    Oerrs Coll
lo0   1536  <Link>                      122910     0   122910     0     0
lo0   1536  127         localhost       122910     0   122910     0     0
tr1   1492  <Link>                      422863     0   309515     0     0
tr1   1492  9.3.4       mickey_sb.itsc. 422863     0   309515     0     0
tr0   1492  <Link>                      229144     0   176978     0     0
tr0   1492  9.3.1       mickey.itsc.aus 229144     0   176978     0     0
```

When the tok0 adapter is unplugged, the service address is swapped to the tok1 adapter.

```
Name  Mtu   Network     Address         Ipkts    Ierrs Opkts    Oerrs Coll
lo0   1536  <Link>                      124881     0   124881     0     0
lo0   1536  127         localhost       124881     0   124881     0     0
tr1   1492  <Link>                         337     0      312     0     0
tr1   1492  9.3.1       mickey.itsc.aus    337     0      312     0     0
tr0   1492  <Link>                           4     0       43    21     0
tr0   1492  9.3.4       mickey_sb.itsc.     4     0       43    21     0
```

At this time, a message appears on the console to warn the users that the node is running in a state where the standby adapter is not available.

```
Adapter 9.3.4.79 is no longer available for use as a standby,
due to either a standby adapter failure or IP address takeover.
```

When the tok1 adapter is reconnected, it is used for the standby address. There is no need for the addresses to swap back to their original adapters. At this time, the following message appeared on the system console:

```
Standby adapter 9.3.4.79 is now available.
```

In all cases, these tests were successful, and we found no differences from the homogeneous case.

9.5.2 Simulated Network Failure

A network failure can be classified into two cases:

local When one node loses all connections to the network.

global When all nodes have lost contact with the network.

We can simulate a local network failure by simultaneously unplugging both of a node's token-ring connections.

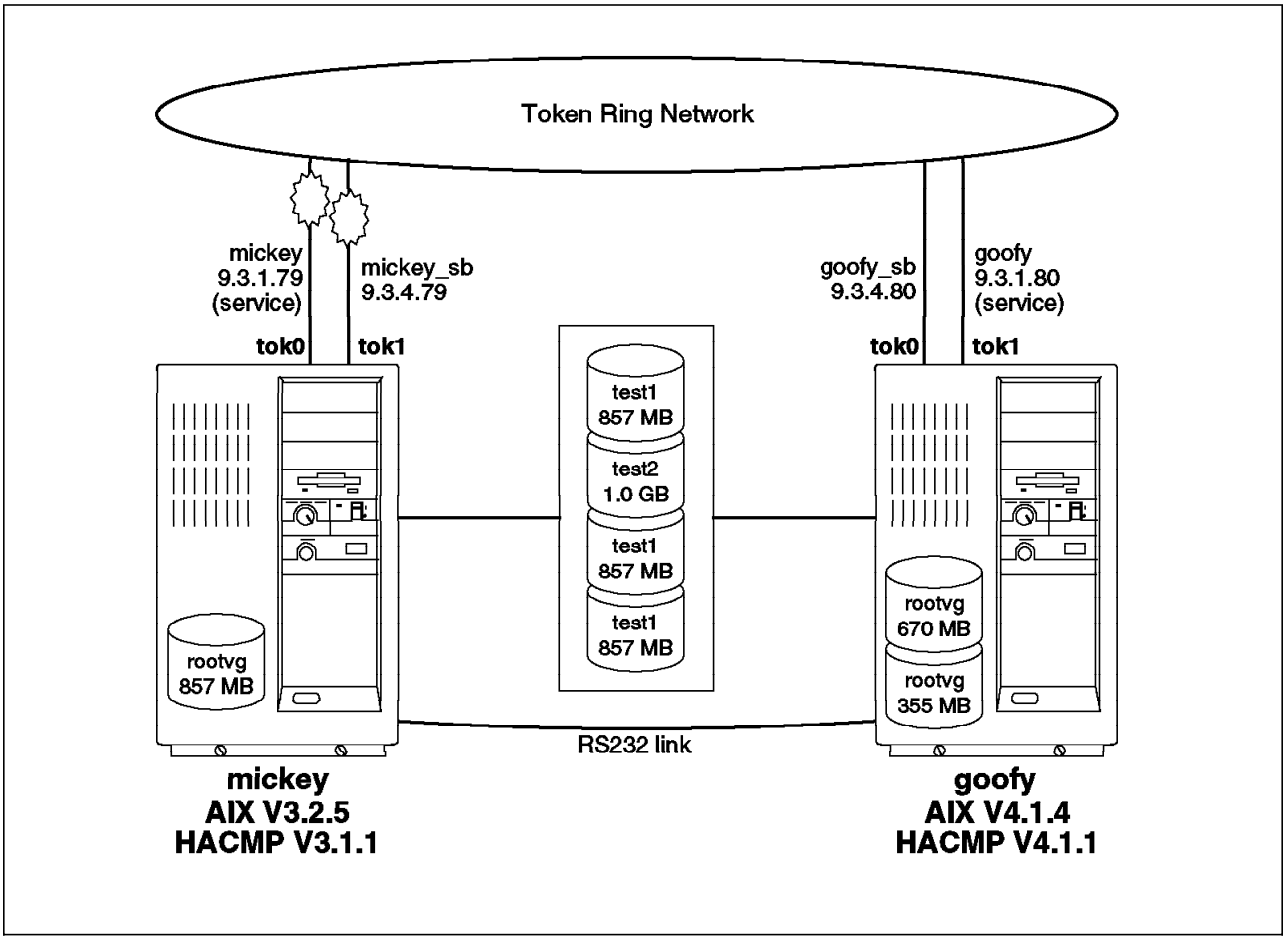

Figure 17. Simulated Local Network Failure

When a network failure is detected, the network_down event script is run; however, by default, this script takes no action. The recovery actions in this situation are highly dependent upon your network configuration and facilities. The event is run, and available to be customized by the user.

The network_down was the least reliable event in our testing; however, this was consistent across the homogeneous and heterogeneous tests. It is very difficult in a two-node cluster for the cluster manager to distinguish between an adapter failure and a local network failure. The output below shows an example of the events that were run, and listed in the /var/adm/cluster.log file, when both the service and standby adapters of node mickey were unplugged from the network.

```
mickey HACMP for AIX: EVENT START: fail_standby 1  9.3.4.79
mickey HACMP for AIX: EVENT COMPLETED: fail_standby 9.3.4.79
mickey HACMP for AIX: EVENT START: network_down 1 trnet1
mickey HACMP for AIX: EVENT COMPLETED: network_down 1 trnet1
mickey HACMP for AIX: EVENT START: network_down_complete 1 trnet1
mickey HACMP for AIX: EVENT COMPLETED: network_down_complete 1 token
```

When the adapters were reconnected, the following events occurred.

```
mickey HACMP for AIX: EVENT START: join_standby 1 9.3.4.79
mickey HACMP for AIX: EVENT COMPLETED: join_standby 1 9.3.4.79
mickey HACMP for AIX: EVENT START: network_up 1 trnet1
mickey HACMP for AIX: EVENT COMPLETED: network_up 1 trnet1
mickey HACMP for AIX: EVENT START: network_up_complete 1 trnet1
mickey HACMP for AIX: EVENT COMPLETED: network_up_complete 1 token 1
```

9.5.3 Complete Node Failure

A complete node failure is induced by powering off one node without first shutting down HACMP or AIX.

In a complete node failure, the service IP address of the failing node is taken over by the standby adapter another node in the cluster. In addition, the cluster can be configured for another node to takeover the disk resources owned by the failing node.

9.5.3.1 IP Address Takeover

This is similar to the situation described in 9.5.1, “Simulated Adapter Failure” on page 259 above; however, the address is taken over by an adapter in another system. This is illustrated in Figure 18 on page 264, Figure 19 on page 265 and Figure 20 on page 266.

The netstat -i output below shows the takeover of the service address from mickey by the standby adapter on goofy when node mickey failed (was switched off). The output before takeover was:

- On mickey:

```
Name  Mtu   Network     Address          Ipkts    Ierrs Opkts    Oerrs Coll
lo0   1536  <Link>                       128486     0   128486     0     0
lo0   1536  127         localhost        128486     0   128486     0     0
tr0   1492  <Link>                         4485     0     3647     0     0
tr0   1492  9.3.1       mickey.itsc.aus    4485     0     3647     0     0
tr1   1492  <Link>                         3131     0     2316     0     0
tr1   1492  9.3.4       mickey_sb.itsc.    3131     0     2316     0     0
```

- On goofy:

```
Name  Mtu   Network     Address            Ipkts Ierrs    Opkts Oerrs  Coll
lo0   16896 <Link>                         1454     0     1653     0     0
lo0   16896 127         localhost          1454     0     1653     0     0
tr1   1492  <Link>10.0.5a.4f.49.42         4048     0     2820     0     0
tr1   1492  9.3.4       goofy_sb.itsc.a    4048     0     2820     0     0
tr0   1492  <Link>10.0.5a.a8.d1.f3         3178     0     2494     0     0
tr0   1492  9.3.1       goofy.itsc.aust    3178     0     2494     0     0
```

When mickey failed, the service address was swapped to the tok1 adapter of goofy:

- On goofy:

```
Name  Mtu   Network     Address            Ipkts Ierrs    Opkts Oerrs  Coll
lo0   16896 <Link>                         2394     0     2600     0     0
lo0   16896 127         localhost          2394     0     2600     0     0
tr0   1492  <Link>10.0.5a.a8.d1.f3         4845     0     3716     0     0
tr0   1492  9.3.1       goofy.itsc.aust    4845     0     3716     0     0
tr1   1492  <Link>10.0.5a.4f.49.42           92     0        4     0     0
tr1   1492  9.3.1       mickey.itsc.aus      92     0        4     0     0
```

The IP address takeover portion of the node failure processing operated with no problems across HACMP versions, regardless of the HACMP Version on the failing node.

9.5.3.2 Disk Resource Takeover

There are two distinct cases when considering disk takeover in the heterogeneous environment. Figure 18 on page 264 shows the state of the cluster before the node failure.

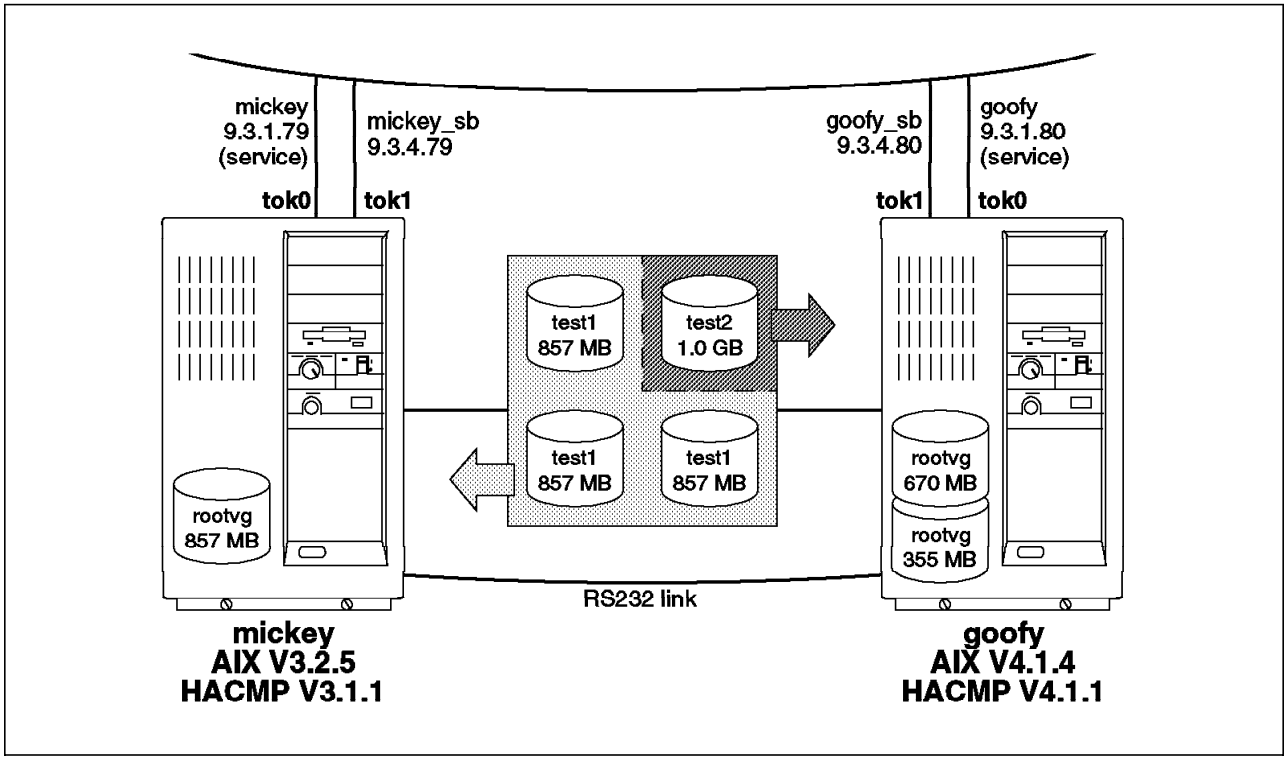

Figure 18. Before Node Failure

Failure of the HACMP Version 3.1.1 Node: This is the successful case. When we switched off node mickey, the event processing proceeded without errors. The state of the cluster after the failure is shown in Figure 19 on page 265.

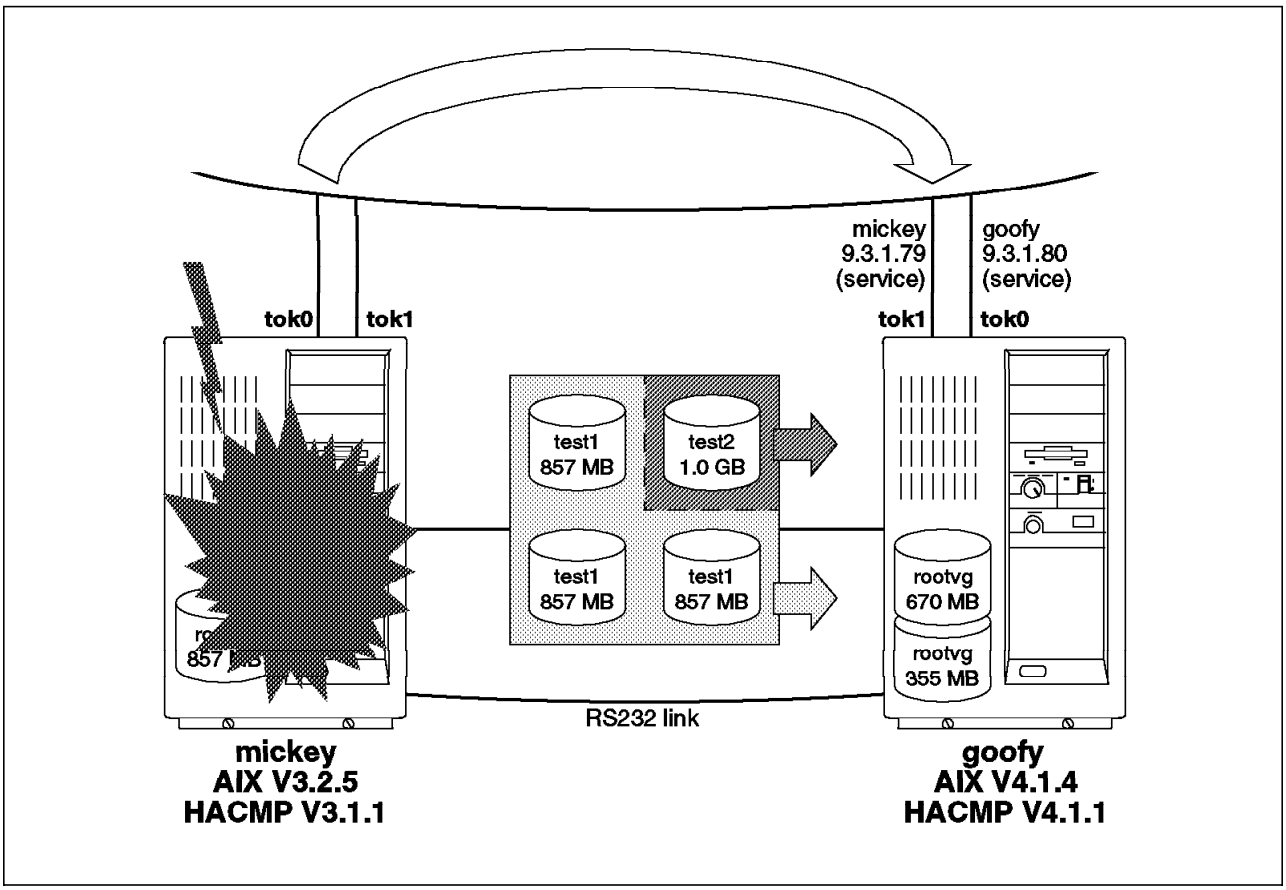

Figure 19. After Failure of HACMP V3.1.1 Node

The output from the df command run on goofy after the failure shows that it has successfully varied on the volume group and mounted the /test1 filesystem.

- On mickey, before the failure:

```
Filesystem    Total KB    free %used   iused %iused Mounted on
/dev/hd4         12288    4824   60%     947    23% /
/dev/hd9var       4096    1828   55%     112    10% /var
/dev/hd2        552960   56916   89%   20260    14% /usr
/dev/hd3         20480   16056   21%      56     0% /tmp
/dev/hd1          4096    3932    4%      17     1% /home
/dev/lvtest1     20480   18560    9%     298     4% /test1
```

- On goofy, before the failure of mickey:

```
Filesystem    512-blocks      Free %Used    Iused %Iused Mounted on
/dev/hd4           24576      6192   75%     1263    31% /
/dev/hd2         1105920    236680   79%    17626    13% /usr
/dev/hd9var        16384     13304   19%      128     7% /var
/dev/hd3           40960     28712   30%       92     2% /tmp
/dev/hd1            8192      7648    7%       38     4% /home
/dev/lvtest2       49152     47552    4%       16     1% /test2
```

Note: The output of the df command has changed from kilobytes in AIX Version 3.2.5 to 512-byte blocks in AIX Version 4.1

- On goofy, after the failure and takeover:

```
Filesystem    512-blocks      Free %Used    Iused %Iused Mounted on
/dev/hd4           24576      6192   75%     1264    31% /
/dev/hd2         1105920    236680   79%    17626    13% /usr
/dev/hd9var        16384     13288   19%      128     7% /var
/dev/hd3           40960     28664   31%       92     2% /tmp
/dev/hd1            8192      7648    7%       38     4% /home
/dev/lvtest2       49152     47552    4%       16     1% /test2
/dev/lvtest1       40960     37120   10%      298     5% /test1
```

Failure of the HACMP Version 4.1.1 Node: As discussed in 8.6.1.3, “Disk Resource Takeover” on page 238, this is the major source of incompatibility between HACMP Version 3.1.1 and Version 4.1.1. Figure 20 shows the state of the cluster after switching off goofy.

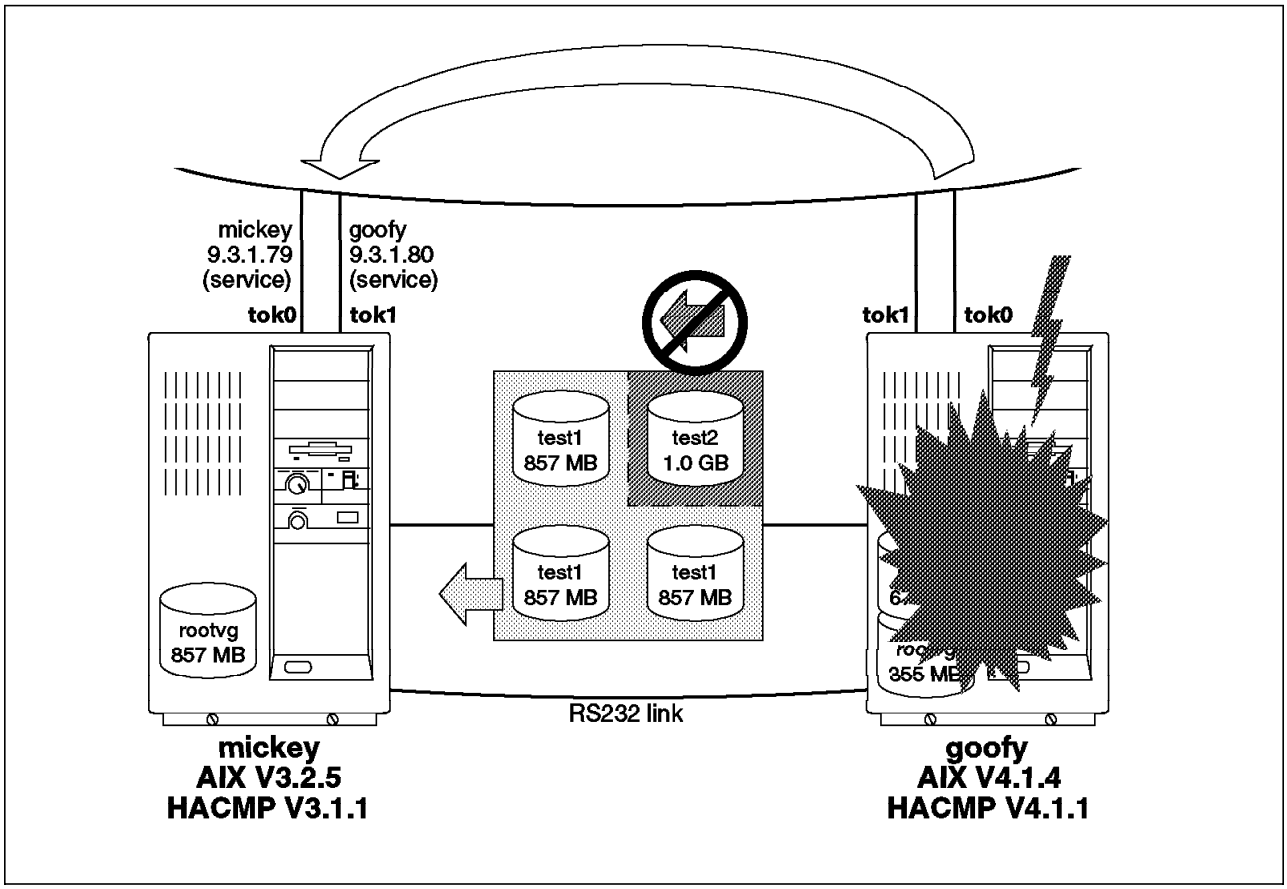

Figure 20. After Failure of HACMP V4.1.1 Node

When node goofy was switched off, we saw the following events listed in the /var/adm/cluster.log file:

```
May 12 14:42:31 mickey HACMP/6000: EVENT START: node_down goofy
May 12 14:42:33 mickey HACMP/6000: EVENT START: node_down_remote goofy
May 12 14:42:34 mickey HACMP/6000: EVENT START: acquire_takeover_addr goofy
May 12 14:42:58 mickey HACMP/6000: EVENT COMPLETED: acquire_takeover_addr goofy
May 12 14:42:58 mickey HACMP/6000: EVENT START: get_disk_vg_fs /test2
May 12 14:43:35 mickey HACMP/6000: /usr/sbin/cluster/events/utils/cl_activate_fs
 : Failed mount of filesystem /test2
May 12 14:43:37 mickey HACMP/6000: EVENT FAILED:1: get_disk_vg_fs /test2
May 12 14:43:42 mickey HACMP/6000: EVENT FAILED:1: node_down_remote goofy
May 12 14:43:43 mickey HACMP/6000: Failure
May 12 14:43:44 mickey HACMP/6000: EVENT FAILED:1: node_down goofy
May 12 14:43:44 mickey clstrmgr[8139]: mickey: bad script status 1 for mickey
May 12 14:43:44 mickey HACMP/6000: EVENT START: event_error mickey node_down goo
fy takeover
May 12 14:43:45 mickey HACMP/6000: EVENT COMPLETED: event_error mickey node_down
 goofy takeover
```

The netstat -i command showed that the IP address had been successfully taken over:

```
Name  Mtu   Network     Address         Ipkts    Ierrs Opkts    Oerrs Coll
lo0   1536  <Link>                      129595     0   129595     0     0
lo0   1536  127         localhost       129595     0   129595     0     0
tr0   1492  <Link>                        1423     0     1176     0     0
tr0   1492  9.3.1       mickey.itsc.aus   1423     0     1176     0     0
tr1   1492  <Link>                         148     0        1     0     0
tr1   1492  9.3.1       goofy.itsc.aust    148     0        1     0     0
```

However, df showed that the filesystem had not been mounted.

After approximately five minutes we saw an additional event:

```
May 12 14:48:31 mickey HACMP/6000: EVENT START: config_too_long 360 node_down go
ofy takeover
```

At this time, the following messages began appearing every 30 seconds on the system console of mickey:

```
Cluster has been in reconfiguration too long.
Cluster has been in reconfiguration too long.
Cluster has been in reconfiguration too long.
Cluster has been in reconfiguration too long.
```

To correct this problem we took the following steps (as described in 8.6.1.3, “Disk Resource Takeover” on page 238).

1. Ran a Recover From Script Failure from SMIT

   This took the node out of the reconfiguration state. The event scripts then completed as shown below:

```
May 12 15:00:54 mickey clstrmgr·8139“: Continue reconfiguration request received.
May 12 15:00:55 mickey HACMP/6000: EVENT START: node_down_complete goofy
May 12 15:00:57 mickey HACMP/6000: EVENT START: node_down_remote_complete goofy
May 12 15:00:58 mickey HACMP/6000: EVENT COMPLETED: node_down_remote_complete go
ofy
May 12 15:00:59 mickey HACMP/6000: EVENT COMPLETED: node_down_complete goofy
```

2. Reinitialized the jfslog

```
# logform /dev/loglvtest2
logform: destroy /dev/loglvtest2 (y)?y
```

3. Checked the filesystem

```
# fsck -y /test2

** Checking /dev/rlvtest2 (/test2
** Phase 0  - Check Log
log redo processing for /dev/rlvtest2
** Phase 1  - Check Blocks and Sizes
** Phase 2  - Check Pathnames
** Phase 3  - Check Connectivity
** Phase 4  - Check Reference Counts
** Phase 5  - Check Inode Map
** Phase 6  - Check Block Map
The superblock is marked dirty. (FIXED)
7 files 1600 blocks 47552 free
***** File system was modified *****
```

4. Mounted the file system

We were then able to see that the filesystem was available for use:

```
#  df
Filesystem    Total KB    free %used   iused %iused Mounted on
/dev/hd4         12288    3760   69%     948    23% /
/dev/hd9var       4096    1820   55%     112    10% /var
/dev/hd2        552960   56912   89%   20260    14% /usr
/dev/hd3         20480   15984   21%      56     0% /tmp
/dev/hd1          4096    3932    4%      17     1% /home
/dev/lvtest1     20480   18560    9%     298     4% /test1
/dev/lvtest2     24576   23776    3%      16     0% /test2
```

9.5.4 Upgrading Node mickey

The procedure used to upgrade mickey was the same as that described above. There is no need to explain each step again.

Once the second upgrade was completed, we again tested various failures in our cluster in a homogeneous HACMP Version 4.1.1 cluster. This time, all events were processed as expected.

Part 5. Appendices

Appendix A. SNA Definitions and Profiles

This appendix lists the SNA definitions and configurations that were used in the sample migration. It includes the VTAM definitions from the mainframe system, as well as the SNA configuration from the systems testcli and testserv.

A.1 VTAM Switched Major Node

The following definitions were made at the host side for both test systems.

```
*********************************************************************    00001000
SWAU001  VBUILD  TYPE=SWNET,                                             +00003000
               MAXGRP=1,                                                 +00004000
               MAXNO=1                                                    00005000
*                                                                         00046000
SC04002  PU    ADDR=01,                                                  +00047000
               IDBLK=071,IDNUM=04002,                                    +00048000
               ANS=CONT,DISCNT=NO,                                       +00049000
               IRETRY=NO,ISTATUS=ACTIVE,                                 +00050000
               MAXDATA=265,MAXOUT=7,                                     +00051000
               MAXPATH=1,                                                +00052000
               PUTYPE=2,SECNET=NO,                                       +00053000
               MODETAB=POKMODE,DLOGMOD=DYNRMT,                           +00054000
               USSTAB=USSRDYN,PACING=1,                                  +00055000
               VPACING=2                                                  00056000
*                                                                         00057000
SC04002A LU    LOCADDR=002,LOGAPPL=SCGVAMP         #TESTSERV             00058000
SC04002B LU    LOCADDR=003                         #TESTSERV             00059000
SC04002C LU    LOCADDR=004                         #TESTCLI              00060000
SC04002D LU    LOCADDR=005                         #TESTCLI              00061000
*
SC04002I LU    LOCADDR=0,DLOGMOD=LU62APPB                                00062000
SC04002J LU    LOCADDR=0,DLOGMOD=LU62APPB                                00063000
SC04002K LU    LOCADDR=0,DLOGMOD=LU62APPA                                00064000
SC04002L LU    LOCADDR=0,DLOGMOD=LU62APPA                                00065000
*                                                                         00066000
```

A.2 Old Profiles on TESTCLI

The following SNA Services/6000 profiles are used on the system TESTCLI before the migration.

Note: The profiles show only those parts that are relevant to the migration. Alias definitions and default profiles are not shown.

```
#SNA  01.02.0101.0315   ***DO NOT MODIFY OR REMOVE***

con201_CONNECTION:
      type = CONNECTION
    profile_name = con201
    attachment_profile_name = att1
    local_lu_profile_name = lu201
    network_name =
    remote_lu_name = lu101
      stop_connection_on_inactivity = no
      lu_type = lu6.2
      interface_type = extended
    remote_tpn_list_name = server
    mode_list_name = server
      node_verification = no
      inactivity_timeout_value = 0
      notify = no
      parallel_sessions = single
      negotiate_session_limits = no
```

```
       security_accepted = none
       conversation_security_access_list_name = CONVDEFAULT


con202_CONNECTION:
       type = CONNECTION
       profile_name = con202
    attachment_profile_name = att1
    local_lu_profile_name = lu202
       network_name =
       remote_lu_name =
       stop_connection_on_inactivity = no
       lu_type = lu2
       interface_type = extended
       remote_tpn_list_name = RDEFAULT
       mode_list_name = MDEFAULT
       node_verification = no
       inactivity_timeout_value = 0
       notify = yes
       parallel_sessions = single
       negotiate_session_limits = no
       security_accepted = none
       conversation_security_access_list_name = CONVDEFAULT


con203_CONNECTION:
       type = CONNECTION
    profile_name = con203
    attachment_profile_name = att1
    local_lu_profile_name = lu203
       network_name =
       remote_lu_name =
       stop_connection_on_inactivity = no
       lu_type = lu2
       interface_type = extended
       remote_tpn_list_name = RDEFAULT
       mode_list_name = MDEFAULT
       node_verification = no
       inactivity_timeout_value = 0
       notify = yes
       parallel_sessions = single
       negotiate_session_limits = no
       security_accepted = none
       conversation_security_access_list_name = CONVDEFAULT


lu201_LOCALLU:
       type = LOCALLU
    profile_name = lu201
    local_lu_name = lu201
    network_name =
       lu_type = lu6.2
       independent_lu = yes
    tpn_list_name = client
       local_lu_address = 1
       sscp_id = *
       number_of_rows = 24
       number_of_columns = 80


lu202_LOCALLU:
       type = LOCALLU
    profile_name = lu202
       local_lu_name =
       network_name =
       lu_type = lu2
       independent_lu = no
       tpn_list_name = TDEFAULT
       local_lu_address = 2
       sscp_id = *
       number_of_rows = 24
       number_of_columns = 80
```

```
lu203_LOCALLU:
       type = LOCALLU
    profile_name = lu203
       local_lu_name =
       network_name =
       lu_type = lu2
       independent_lu = no
       tpn_list_name = TDEFAULT
       local_lu_address = 3
       sscp_id = *
       number_of_rows = 24
       number_of_columns = 80


att1_ATTACHMENT:
       type = ATTACHMENT
       profile_name = att1
       control_point_profile_name = CP2
       logical_link_profile_name = con2ll
       physical_link_profile_name = con2pl
       logical_link_type = token_ring
       restart_on_deactivation = no
       stop_attachment_on_inactivity = no
       station_type = secondary
       physical_link_type = token_ring
       remote_secondary_station_address = 1
       smart_modem_command_sequence =
       length_of_command_sequence = 0
       call_type = call
       x25_level = 1984
       listen_name = IBMQLLC
       autolisten = no
       timeout_value = 0
       remote_link_name_ethernet =
       remote_link_name_token_ring =
       remote_link_address = 400052004002
       selection_sequence =
       length_of_selection_sequence = 0
       network_type = switched
       access_routing = link_address
       remote_sap_address = 04
       remote_sap_address_range_lower = 04
       remote_sap_address_range_upper = EC
       virtual_circuit_type = permanent
       remote_station_X.25_address =
       optional_X.25_facilities = no
       logical_channel_number_of_PVC = 1
       reverse_charging = no
       rpoa = no
       default_packet_size = no
       default_window_size = no
       default_throughput_class = no
       closed_user_group = no
       closed_user_group_outgoing = no
       network_user_id = no
       network_user_id_name =
       data_network_identification_code =
       packet_size_for_received_data = 128
       packet_size_for_transmit_data = 128
       window_size_for_received_data = 2
       window_size_for_transmit_data = 2
       throughput_class_for_received_data = 9600
       throughput_class_for_transmit_data = 9600
       index_to_selected_closed_user_group = 0
       lu_address_registration = no
       lu_address_registration_name = LDEFAULT


SDEFAULT_SNA:
       type = SNA
       profile_name = SDEFAULT
       total_active_open_connections = 200
       total_sessions = 200
```

```
       total_conversations = 200
       server_synonym_name =
       nmvt_action_when_no_nmvt_process = reject
       restart_action =
       stdin =
       stdout =
       stderr =
       sna_error_log = no


servrcv_REMOTETPN:
       type = REMOTETPN
    profile_name = servrcv
    tpn_name = CPICRCV
    tpn_name_hex =
       pip_data = no
       conversation_type = basic
       recovery_level = no_reconnect
       sync_level = confirm
       tpn_name_in_hex = no


servmsg_REMOTETPN:
       type = REMOTETPN
    profile_name = servmsg
    tpn_name = SUBRRECV
    tpn_name_hex =
       pip_data = no
       conversation_type = basic
       recovery_level = no_reconnect
       sync_level = confirm
       tpn_name_in_hex = no


server_REMOTETPNLIST:
       type = REMOTETPNLIST
    Listname = server
    list_members = servrcv
    list_members = servmsg


REMOTE_TPN_REMOTETPNLIST:
       type = REMOTETPNLIST
       Listname = REMOTE_TPN
       list_members = RDEFAULT


receive_TPN:
       type = TPN
    profile_name = receive
    tpn_name = CPICRCV
    tpn_name_hex =
       conversation_type = mapped
       pip_data = no
       sync_level = either
       recovery_level = no_reconnect
    full_path_to_tpn_executable = /usr/lpp/sna/sample/bin/CPICRCV
      multiple_instances = yes
       user_id = 404
    server_synonym_name = client
       restart_action = once
       communication_type = signals
       stdin = /dev/null
       stdout = /dev/console
       stderr = /dev/console
       subfields = 0
       communication_ipc_queue_key = 0
       tpn_name_in_hex = no
       security_required = none
       resource_security_access_list_name = RSRCDEFAULT
```

```
message_TPN:
       type = TPN
    profile_name = message
    tpn_name = SUBRRECV
    tpn_name_hex =
       conversation_type = mapped
       pip_data = no
       sync_level = either
       recovery_level = no_reconnect
    full_path_to_tpn_executable = /usr/lpp/sna/sample/bin/SUBRRECV
       multiple_instances = yes
       user_id = 404
    server_synonym_name = client
       restart_action = once
       communication_type = signals
       stdin = /dev/null
       stdout = /dev/console
       stderr = /dev/console
       subfields = 0
       communication_ipc_queue_key = 0
       tpn_name_in_hex = no
       security_required = none
       resource_security_access_list_name = RSRCDEFAULT


client_TPNLIST:
       type = TPNLIST
    Listname = client
    list_members = receive
    list_members = message


LOCAL_TPN_TPNLIST:
       type = TPNLIST
       Listname = LOCAL_TPN
       list_members = TDEFAULT


CDEFAULT_CONTROLPOINT:
       type = CONTROLPOINT
       profile_name = CDEFAULT
       xid_node_id = 07100000
       network_name =
       cp_name =



CP2_CONTROLPOINT:
       type = CONTROLPOINT
    profile_name = CP2
    xid_node_id = 071D0500
       network_name =
       cp_name =



client_MODELIST:
       type = MODELIST
    Listname = client
    list_members = cpic

cpic_MODE:
       type = MODE
    profile_name = cpic
       mode_name = MDEFAULT
       maximum_number_of_sessions = 8
       minimum_contention_winners = 4
       minimum_contention_losers = 0
       receive_pacing = 32
       send_pacing = 32
       maximum_ru_size = 3840
       recovery_level = no_reconnect


con2ll_TOKENRINGLOGICAL:
```

```
        type = TOKENRINGLOGICAL
        profile_name = con2ll
        retry_limit = 20
        transmit_window_count = 10
        dynamic_window_increment = 1
        retransmit_count = 8
        receive_window_count = 127
        ring_access_priority = 0
        inactivity_timeout = 48
        drop_link_on_inactivity = yes
        response_timeout = 2
        acknowledgement_timeout = 1
        force_disconnect_timeout = 120
        link_trace = no
        trace_entry_size = long
        logical_link_type = token_ring
        maximum_i_field = system_defined
        maximum_i_field_size = 30729
        physical_link_type = token_ring


con2pl_TOKENRINGPHYSICAL:
        type = TOKENRINGPHYSICAL
        profile_name = con2pl
        device_name = tok0
        local_link_name =
        local_sap_address = 04
        physical_link_type = token_ring
        maximum_number_of_logical_links = 32
```

A.3 Migrated Profiles on TESTCLI

These profiles are result of using the migratesna command to migrate the previous profiles.

Note: Most default profiles are not shown.

```
link_station_token_ring:
        prof_name                       = "att1"
        sna_dlc_profile_name            = "att1"
        LU_registration_supported       = no
        LU_registration_profile_name    = ""
        cp_cp_sessions_supported        = no
        link_tracing                    = no
        trace_format                    = long
        xid_node_id                     = 071D0500
        use_control_pt_xid              = no
        call_out_on_activation          = yes
        restart_on_activation           = no
        remote_link_name                = ""
        remote_link_address             = 400052004002
        access_routing_type             = link_address
        remote_sap                      = 04


control_pt:
        prof_name                       = "node_cp"
        network_name                    = "NET1"
        control_pt_name                 = "CP2"


session_lu2:
        prof_name                       = "con202"
        link_station_profile_name       = "att1"
        local_lu_name                   = ""
        local_lu_address                = 2
        sscp_id                         = "*"
        network_name                    = ""
        remote_lu_name                  = ""
        max_rows                        = 24
        max_columns                     = 80
```

```
session_lu2:
        prof_name                       = "con203"
        link_station_profile_name       = "att1"
        local_lu_name                   = ""
        local_lu_address                = 3
        sscp_id                         = "*"
        network_name                    = ""
        remote_lu_name                  = ""
        max_rows                        = 24
        max_columns                     = 80


side_info:
        prof_name                       = "IDEFAULT"
        remote_tp_name                  = "rtpn"
        remote_tp_name_in_hex           = no


side_info:
        prof_name                       = "con201"
        local_lu_or_control_pt_alias    = "LU201"
        partner_lu_alias                = "LU101"
        remote_tp_name                  = "CPICRCV"
        remote_tp_name_in_hex           = no


local_lu_lu6.2:
        prof_name                       = "LU62DEFAULT"
        local_lu_name                   = ""
        local_lu_address                = 1
        sscp_id                         = "*"
        local_lu_alias                  = ""
        local_lu_dependent              = no
        conversation_security_list_profile_name = ""


local_lu_lu6.2:
        prof_name                       = "lu201"
        local_lu_name                   = "LU201"
        local_lu_address                = 1
        link_station_prof_name          = "att1"
        sscp_id                         = "*"
        local_lu_alias                  = "LU201"
        local_lu_dependent              = no
        conversation_security_list_profile_name = ""


local_tp:
        prof_name                       = "receive"
        tp_name                         = "CPICRCV"
        conversation_type               = mapped
        pip_data_present                = no
        sync_level                      = none/confirm
        full_path_tp_exe                = "/usr/lpp/sna/sample/bin/CPICRCV"
        multiple_instances              = yes
        user_id                         = 404
        server_synonym_name             = "client"
        restart_action                  = once
        communication_type              = signals
        standard_input_device           = "/dev/null"
        standard_output_device          = "/dev/console"
        standard_error_device           = "/dev/console"
        pip_data_subfields_number       = 0
        ipc_queue_key                   = 0
        tp_name_in_hex                  = no
        resource_security_level         = none
        resource_access_list_profile_name   = ""


local_tp:
        prof_name                       = "message"
        tp_name                         = "SUBRRECV"
        conversation_type               = mapped
        pip_data_present                = no
```

```
        sync_level                      = none/confirm
        full_path_tp_exe                = "/usr/lpp/sna/sample/bin/SUBRRECV"
        multiple_instances              = yes
        user_id                         = 404
        server_synonym_name             = "client"
        restart_action                  = once
        communication_type              = signals
        standard_input_device           = "/dev/null"
        standard_output_device          = "/dev/console"
        standard_error_device           = "/dev/console"
        pip_data_subfields_number       = 0
        ipc_queue_key                   = 0
        tp_name_in_hex                  = no
        resource_security_level         = none
        resource_access_list_profile_name   = ""


lu_reg:
        prof_name                       = "LUREGDEFAULT"
        lu_address_registered_list      = {64}


rsc_list:
        prof_name                       = "RSCDEFAULT"
        username_list                   = {}


conv_list:
        prof_name                       = "CONVDEFAULT"
        username_list                   = {}



mode:
        prof_name                       = "cpic"
        mode_name                       = "MDEFAULT"
        max_sessions                    = 8
        min_conwinner_sessions          = 4
        min_conloser_sessions           = 0
        receive_pacing_window           = 32
        max_ru_size                     = 3840


side_info:
        prof_name                       = "IDEFAULT"
        remote_tp_name                  = ""
        mode_name                       = ""


sna_dlc_token_ring:
        prof_name                       = "TLINKDEFAULT"
        datalink_device_name            = "tok0"
        force_timeout                   = 120
        user_defined_max_i_field        = no
        max_i_field_length              = 30729
        retry_limit                     = 20
        link_name                       = ""
        local_sap                       = 04
        max_active_link_stations        = 32
        transmit_window_count           = 127
        dynamic_window_increment        = 1
        retransmit_count                = 8
        receive_window_count            = 1
        priority                        = 0
        inact_timeout                   = 48
        response_timeout                = 4
        acknowledgement_timeout         = 1


sna_dlc_token_ring:
        prof_name                       = "att1"
        datalink_device_name            = "tok0"
        force_timeout                   = 120
        user_defined_max_i_field        = no
        max_i_field_length              = 30729
```

```
        retry_limit                     = 20
        link_name                       = ""
        local_sap                       = 04
        max_active_link_stations        = 32
        transmit_window_count           = 127
        dynamic_window_increment        = 1
        retransmit_count                = 8
        receive_window_count            = 1
        priority                        = 0
        inact_timeout                   = 48
        response_timeout                = 4
        acknowledgement_timeout         = 1


partner_lu6.2:
        prof_name                       = "PLUDEFAULT"
        parallel_session_supp           = no
        conversation_security_level     = none
        session_security_supp           = no


partner_lu6.2:
        prof_name                       = "con201"
        partner_lu_alias                = "LU101"
        parallel_session_supp           = no
        conversation_security_level     = none
        session_security_supp           = no
```

A.4 SNA Profiles on TESTSERV

The following profiles were used on the system TESTSERV. Since this system was initially running SNA Server Version 2.1, no modifications to the profiles were required during the migration.

```
sna:
       prof_name                                    = "sna"
       max_sessions                                 = 200
       max_conversations                            = 200
       restart_action                               = once
       dynamic_inbound_partner_lu_definitions_allowed = yes
       standard_output_device                       = "/dev/console"
       standard_error_device                        = "/var/sna/sna.stderr"
       nmvt_action_when_no_nmvt_process             = reject
       trusted_group_ids                            = {system}
       sense_detail_level                           = specific
       start_snmp_subagent                          = no
       limited_resource_timeout                     = no
       limited_resource_timeout_value               = 15
       comments                                     = ""

control_pt:
       prof_name                                    = "node_cp"
       xid_node_id                                  = "*"
       network_name                                 = "NET1"
       control_pt_name_alias                        = "CP1"
       control_pt_name                              = "CP1"
       control_pt_node_type                         = appn_end_node
       max_cached_trees                             = 500
       max_nodes_in_topology_database               = 500
       route_addition_resistance                    = 128
       comments                                     = ""

session_lu2:
       prof_name                                    = "toHOST02"
       link_station_profile_name                    = "toHOST"
       local_lu_name                                = ""
       network_name                                 = ""
       remote_lu_name                               = ""
       comments                                     = ""
       local_lu_address                             = 2
       sscp_id                                      = *
       max_rows                                     = 24
       max_columns                                  = 80
```

```
       comments                                        = ""

session_lu2:
       prof_name                                       = "toHOST03"
       link_station_profile_name                       = "toHOST"
       local_lu_name                                   = ""
       network_name                                    = ""
       remote_lu_name                                  = ""
       comments                                        = ""
       local_lu_address                                = 3
       sscp_id                                         = *
       max_rows                                        = 24
       max_columns                                     = 80
       comments                                        = ""

local_lu_lu6.2:
       prof_name                                       = "LU101"
       local_lu_name                                   = "LU101"
       local_lu_alias                                  = "LU101"
       local_lu_dependent                              = no
       local_lu_address                                =
       sscp_id                                         = *
       link_station_prof_name                          = "toTESTCL"
       conversation_security_list_profile_name         = ""
       rrm_enabled                                     = no
       comments                                        = ""

partner_lu6.2:
       prof_name                                       = "lu201"
       fq_partner_lu_name                              = "NET1.LU201"
       partner_lu_alias                                = "LU201"
       session_security_supp                           = no
       parallel_session_supp                           = yes
       conversation_security_level                     = none
       comments                                        = ""

partner_lu6.2_location:
       prof_name                                       = "lu201"
       fq_partner_lu_name                              = "NET1.LU201"
       partner_location_method                         = link_station
       fq_partner_owning_cp_name                       = ""
       local_node_is_network_server_for_len_node       = no
       fq_node_server_name                             = ""
       local_lu_name                                   = "LU101"
       link_station_profile_name                       = "toTESTCL"
       comments                                        = ""

side_info:
       prof_name                                       = "lu201"
       local_lu_or_control_pt_alias                    = "LU101"
       partner_lu_alias                                = "LU201"
       fq_partner_lu_name                              = ""
       mode_name                                       = ""
       remote_tp_name_in_hex                           = no
       remote_tp_name                                  = "CPICRCV"
       comments                                        = ""

local_tp:
       prof_name                                       = "receive"
       tp_name                                         = "CPICRCV"
       tp_name_in_hex                                  = no
       pip_data_present                                = no
       pip_data_subfields_number                       = 0
       command_line_parameters_present                 = no
       command_line_parameters                         = ""
       conversation_type                               = mapped
       sync_level                                      = none/confirm
       resource_security_level                         = none
       resource_access_list_profile_name               = ""
       full_path_tp_exe                                = "/usr/lpp/sna/samples/bin/CPICRCV"
       multiple_instances                              = no
       user_id                                         = 0
       server_synonym_name                             = ""
       restart_action                                  = once
       communication_type                              = signals
```

```
        ipc_queue_key                           = 0
        attach_timeout                          = yes
        attach_timeout_value                    = 60
        standard_input_device                   = "/dev/console"
        standard_output_device                  = "/dev/console"
        standard_error_device                   = "/dev/console"
        comments                                = ""

local_tp:
        prof_name                               = "message"
        tp_name                                 = "SUBRRECV"
        tp_name_in_hex                          = no
        pip_data_present                        = no
        pip_data_subfields_number               = 0
        command_line_parameters_present         = no
        command_line_parameters                 = ""
        conversation_type                       = mapped
        sync_level                              = none/confirm
        resource_security_level                 = none
        resource_access_list_profile_name       = ""
        full_path_tp_exe                        = "/usr/lpp/sna/samples/bin/SUBRRECV"
        multiple_instances                      = no
        user_id                                 = 0
        server_synonym_name                     = ""
        restart_action                          = once
        communication_type                      = signals
        ipc_queue_key                           = 0
        attach_timeout                          = yes
        attach_timeout_value                    = 60
        standard_input_device                   = "/dev/console"
        standard_output_device                  = "/dev/console"
        standard_error_device                   = "/dev/console"
        comments                                = ""

link_station_token_ring:
        prof_name                               = "toHOST"
        use_control_pt_xid                      = no
        xid_node_id                             = 0x07104002
        sna_dlc_profile_name                    = "toHOST"
        stop_on_inactivity                      = no
        time_out_value                          = 0
        LU_registration_supported               = no
        LU_registration_profile_name            = ""
        link_tracing                            = no
        trace_format                            = long
        access_routing_type                     = link_address
        remote_link_name                        = ""
        remote_link_address                     = 0x400008210210
        remote_sap                              = 0x04
        call_out_on_activation                  = yes
        verify_adjacent_node                    = no
        net_id_of_adjacent_node                 = ""
        cp_name_of_adjacent_node                = ""
        xid_node_id_of_adjacent_node            = "*"
        node_type_of_adjacent_node              = learn
        solicit_sscp_sessions                   = yes
        activate_link_during_system_init        = no
        activate_link_on_demand                 = no
        cp_cp_sessions_supported                = yes
        cp_cp_session_support_required          = no
        adjacent_node_is_preferred_server       = no
        initial_tg_number                       = 0
        restart_on_normal_deactivation          = no
        restart_on_abnormal_deactivation        = no
        restart_on_activation                   = no
        TG_effective_capacity                   = 4300800
        TG_connect_cost_per_time                = 0
        TG_cost_per_byte                        = 0
        TG_security                             = nonsecure
        TG_propagation_delay                    = lan
        TG_user_defined_1                       = 128
        TG_user_defined_2                       = 128
        TG_user_defined_3                       = 128
        comments                                = ""
```

```
link_station_token_ring:
       prof_name                                = "toTESTCL"
       use_control_pt_xid                       = yes
       xid_node_id                              = "*"
       sna_dlc_profile_name                     = "toHOST"
       stop_on_inactivity                       = no
       time_out_value                           = 0
       LU_registration_supported                = no
       LU_registration_profile_name             = ""
       link_tracing                             = no
       trace_format                             = long
       access_routing_type                      = link_address
       remote_link_name                         = ""
       remote_link_address                      = ""
       remote_sap                               = 0x04
       call_out_on_activation                   = no
       verify_adjacent_node                     = no
       net_id_of_adjacent_node                  = ""
       cp_name_of_adjacent_node                 = ""
       xid_node_id_of_adjacent_node             = "*"
       node_type_of_adjacent_node               = learn
       solicit_sscp_sessions                    = yes
       activate_link_during_system_init         = yes
       activate_link_on_demand                  = no
       cp_cp_sessions_supported                 = yes
       cp_cp_session_support_required           = no
       adjacent_node_is_preferred_server        = no
       initial_tg_number                        = 0
       restart_on_normal_deactivation           = yes
       restart_on_abnormal_deactivation         = yes
       restart_on_activation                    = no
       TG_effective_capacity                    = 4300800
       TG_connect_cost_per_time                 = 0
       TG_cost_per_byte                         = 0
       TG_security                              = nonsecure
       TG_propagation_delay                     = lan
       TG_user_defined_1                        = 128
       TG_user_defined_2                        = 128
       TG_user_defined_3                        = 128
       comments                                 = ""

sna_dlc_token_ring:
       prof_name                                = "toHOST"
       datalink_device_name                     = "tok0"
       force_timeout                            = 120
       user_defined_max_i_field                 = no
       max_i_field_length                       = 30729
       max_active_link_stations                 = 100
       num_reserved_inbound_activation          = 0
       num_reserved_outbound_activation         = 0
       transmit_window_count                    = 127
       dynamic_window_increment                 = 1
       retransmit_count                         = 8
       receive_window_count                     = 1
       priority                                 = 0
       inact_timeout                            = 48
       response_timeout                         = 4
       acknowledgement_timeout                  = 1
       link_name                                = ""
       local_sap                                = 0x04
       retry_interval                           = 60
       retry_limit                              = 20
       dynamic_link_station_supported           = yes
       trace_base_listen_link_station           = no
       trace_base_listen_link_station_format    = long
       dynamic_lnk_solicit_sscp_sessions        = yes
       dynamic_lnk_cp_cp_sessions_supported     = yes
       dynamic_lnk_cp_cp_session_support_required  = no
       dynamic_lnk_TG_effective_capacity        = 4300800
       dynamic_lnk_TG_connect_cost_per_time     = 0
       dynamic_lnk_TG_cost_per_byte             = 0
       dynamic_lnk_TG_security                  = nonsecure
       dynamic_lnk_TG_propagation_delay         = lan
       dynamic_lnk_TG_user_defined_1            = 128
       dynamic_lnk_TG_user_defined_2            = 128
```

```
       dynamic_lnk_TG_user_defined_3                = 128
       comments                                     = ""

mode:
       prof_name                                    = "DFLTMODE"
       mode_name                                    = "DFLTMODE"
       max_sessions                                 = 8
       min_conwinner_sessions                       = 4
       min_conloser_sessions                        = 0
       auto_activate_limit                          = 0
       max_adaptive_receive_pacing_window           = 16
       receive_pacing_window                        = 7
       max_ru_size                                  = 1024
       min_ru_size                                  = 256
       class_of_service_name                        = "#CONNECT"
       comments                                     = ""

mode:
       prof_name                                    = "SNACKETS"
       mode_name                                    = "SNACKETS"
       max_sessions                                 = 100
       min_conwinner_sessions                       = 50
       min_conloser_sessions                        = 0
       auto_activate_limit                          = 0
       max_adaptive_receive_pacing_window           = 16
       receive_pacing_window                        = 7
       max_ru_size                                  = 3840
       min_ru_size                                  = 128
       class_of_service_name                        = "#CONNECT"
       comments                                     = ""

mode:
       prof_name                                    = "cpic"
       mode_name                                    = "MDEFAULT"
       max_sessions                                 = 8
       min_conwinner_sessions                       = 4
       min_conloser_sessions                        = 0
       auto_activate_limit                          = 0
       max_adaptive_receive_pacing_window           = 16
       receive_pacing_window                        = 32
       max_ru_size                                  = 3840
       min_ru_size                                  = 256
       class_of_service_name                        = "#CONNECT"
       comments                                     = ""

dwnstrm_lu:
       prof_name                                    = "TSTCLI02"
       downstream_link_station_prof_name            = "toTESTCL"
       downstream_PU_name                           = "CP2PU"
       downstream_lu_address                        = 2
       auto_logoff                                  = yes
       gateway_host_prof_name                       = "hostlu4"
       comments                                     = ""

dwnstrm_lu:
       prof_name                                    = "TSTCLI03"
       downstream_link_station_prof_name            = "toTESTCL"
       downstream_PU_name                           = "CP2PU"
       downstream_lu_address                        = 3
       auto_logoff                                  = yes
       gateway_host_prof_name                       = "hostlu5"
       comments                                     = ""

glhostgrp:
       prof_name                                    = "hostlu4"
       lu_type                                      = lu2
       comments                                     = ""
       gateway_host_definition                      1
          ghost_def_prof_name                       = "tohost"
          lu_pooling                                = dedicated
          pool_class_name                           = ""
          host_lu_address                           = 4
          inactivity_timeout                        = 0

glhostgrp:
```

```
       prof_name                                = "hostlu5"
       lu_type                                  = lu2
       comments                                 = ""
       gateway_host_definition                  1
          ghost_def_prof_name                   = "tohost"
          lu_pooling                            = dedicated
          pool_class_name                       = ""
          host_lu_address                       = 5
          inactivity_timeout                    = 0

ghostdef:
       prof_name                                = "tohost"
       host_link_station_prof_name              = "toHOST"
       minimize_link_usage                      = yes
       comments                                 = ""

mptn_env:
       prof_name                                = "env_values"
       sna_suffix                               = "SNA.IBM.COM"
       connection_retry                         = 300
       connection_wait                          = 30
       inactivity_timer                         = 120
       mptn_well_known_port                     = 397
       unacknowledged_dg_retry                  = 30
       unsent_dg_retry                          = 3
       comments                                 = ""
```

List of Abbreviations

AFS	Andrew File System
APA	All Points Addressable
API	Application Programming Interface
APPC	Advanced Program-to-Program Communications
APPN	Advanced Peer-to-Peer Networking
ATE	Asynchronous Terminal Emulation
BFF	Backup File Format
BLKMUX	Block Multiplexer Channel
BMPX/BMX	Block Multiplexer Channel
BOS	Base Operating System
BSD	Berkeley Software Distribution
CLIOS	Client Input Output/Sockets
COMIO	Common Input/Output
CP	Control Point
CPI-C	Common Programming Interface for Communications
DLC	Data Link Control
DLPI	Data Link Programming Interface
DWM	Diskless Workstation Management
EN	End Node
EOF	End of File
ESCON	Enterprise Systems Connection
FDDI	Fiber Distributed Data Interface
FTP	File Transfer Protocol
HCON	3270 Host Connection Program/6000
HLLAPI	High-Level Language Application Programming Interface
IBM	International Business Machines Corporation
ISA	Industry Standard Architecture
ITSO	International Technical Support Organization
LAN	Local Area Network
LEN	Low-Entry Network
LFT	Low-Function Terminal
LPP	Licensed Program Product
LU	Logical Unit
NAU	Network Accessible Units
NCP	Network Control Program
NFS	Network File System
NIM	Network Installation Management
NIS	Network Information System
NN	Network Node
NPI	Network Provider Interface
NUA	Net Usable Area
ODM	Object Data Manager
OPP	Optional Program Products
PAD	Packet Assembler-Disassembler
PCI	Peripheral Component Interconnect
PROFS	Professional Office System
PP	Physical Partition
PU	Physical Unit
RDBMS	Relational Database Management System
RJE	Remote Job Entry
RTPC	Remote Transaction Program Conversation
RTPN	Remote Transaction Program Name
SAP	Service Access Point
SDLC	Synchronous Data Link Control
SIPO	System Installation Productivity Option
SNA	System Network Architecture
SNMP	Simple Network Management Protocol
SMP	Symmetric Multiprocessor
SOA	Start of Authority
SPO	System Program Offering (stacked Tape for AIX products)
SRC	System Resource Controller

SSCP	System Services Control Program
PTF	Program Temporary Fix
TCB	Trusted Computing Base
TOC	Table of Contents
TP	Transaction Program
TPN	Transaction Program Name
VRMF	Version Release Modification Fix
VTAM	Virtual Telecommunications Access Method
XID	Exchange Identifier

Index

Special Characters

Numerics

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

Q

R

S

T

U

V

W

X

ITSO Redbook Evaluation

International Technical Support Organization
A Holistic Approach to AIX V4.1 Migration, Volume 2
TCP/IP, SNA, HACMP and Multiple Systems
May 1996

Publication No. SG24-4653-00

Your feedback is very important to help us maintain the quality of ITSO redbooks. **Please fill out this questionnaire and return it using one of the following methods:**

- Mail it to the address on the back (postage paid in U.S. only)
- Give it to an IBM marketing representative for mailing
- Fax it to: Your International Access Code + 1 914 432 8246
- Send a note to REDBOOK@VNET.IBM.COM

Please rate on a scale of 1 to 5 the subjects below.
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction	____		
Organization of the book	____	Grammar/punctuation/spelling	____
Accuracy of the information	____	Ease of reading and understanding	____
Relevance of the information	____	Ease of finding information	____
Completeness of the information	____	Level of technical detail	____
Value of illustrations	____	Print quality	____

Please answer the following questions:

a) Are you an employee of IBM or its subsidiaries: Yes____ No____

b) Do you work in the USA? Yes____ No____

c) Was this redbook published in time for your needs? Yes____ No____

d) Did this redbook meet your needs? Yes____ No____

If no, please explain:

_____

_____

What other topics would you like to see in this redbook?

_____

_____

What other redbooks would you like to see published?

_____

Comments/Suggestions: **(THANK YOU FOR YOUR FEEDBACK!)**

Name	Address
Company or Organization	
Phone No.	

ITSO Redbook Evaluation
SG24-4653-00

Cut or Fold
Along Line

Fold and Tape **Please do not staple** Fold and Tape

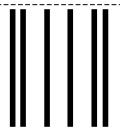

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

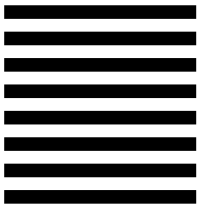

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM International Technical Support Organization
Department HZ8, Building 678
P.O. BOX 12195
RESEARCH TRIANGLE PARK NC
USA 27709-2195

Fold and Tape **Please do not staple** Fold and Tape

Cut or Fold
Along Line

Printed in U.S.A.