International Technical Support Organization

SG24-4553-00

An HACMP Cookbook

December 1995

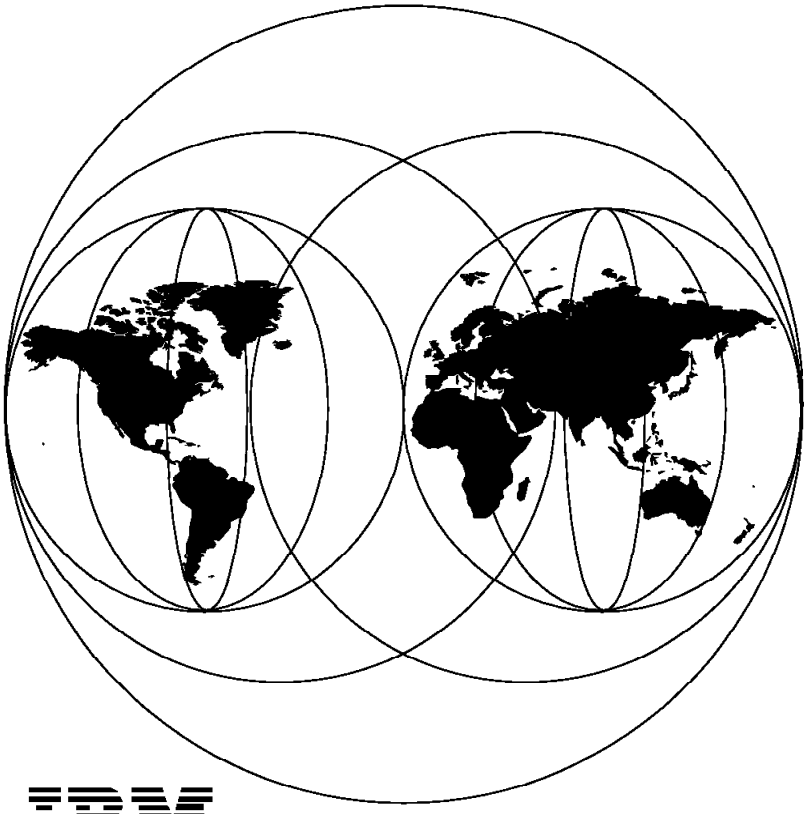

IBM

International Technical Support Organization

Austin Center

International Technical Support Organization

SG24-4553-00

An HACMP Cookbook

December 1995

Take Note!

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

First Edition (December 1995)

This edition applies to Version 3.1.1 of HACMP/6000, Program Number 5696-923, for use with the AIX/6000 Operating System Version 3.2.5.

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 045 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

This document deals with HACMP/6000 Version 3.1.1. Its goal is to serve as a reminder, checklist and operating guide for the steps required in order to install and customize HACMP/6000.

It describes a set of tools developed by the HACMP services team in IBM France, which make it easier to design, customize and document an HACMP cluster. Included in the book are the following:

- How to install the HACMP product
- Description of the tools developed by the HACMP services team in IBM France
- Steps to be carried out during an installation, including customization
- Testing suggestions

Following the instructions in the checklist will assist you towards a smooth and error-free installation. A basic understanding of the HACMP is assumed, and therefore is not included in the book.

(215 pages)

Contents

Figures

Tables

Special Notices

This publication is intended to help customers and IBM services personnel to more easily plan, install, set up, and document their HACMP clusters. The information in this publication is not intended as the specification of any programming interfaces that are provided by HACMP/6000 Version 3.1.1. See the PUBLICATIONS section of the IBM Programming Announcement for HACMP Version 3.1.1 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 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.

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:

AIX	HACMP/6000
IBM	OS/2
POWERserver	POWERstation
RISC System/6000	RS/6000
SP	

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc.

Preface

This publication is intended to help customers and IBM services personnel to more easily plan, install, set up, and document their HACMP clusters. It contains a description of a set of tools developed by the professional services team of IBM France for this purpose.

This document is intended for anyone who needs to implement an HACMP cluster.

How This Document is Organized

The document is organized as follows:

- Chapter 1, “Overview of the Tools”

  This chapter briefly describes each of the configuration and documentation tools included with the book.

- Chapter 2, “Inventory Tool”

  This chapter includes a description of and sample output from a tool that takes an initial inventory of a system that will be a cluster node, and reports any potential problems.

- Chapter 3, “Setting up a Cluster”

  This chapter begins the description of setting up our example cluster. It introduces and describes the example cluster we will set up and use throughout the book, and covers the major planning considerations to be made before starting a cluster setup.

- Chapter 4, “Pre-Installation Activities”

  The set of AIX configuration tasks that need to be done before the installation of HACMP is covered in this chapter. This includes TCP/IP network adapter definitions, tty and SCSI target mode definitions, connecting shared disks, and defining shared volume groups.

- Chapter 5, “Installing the HACMP/6000 Software”

  This chapter describes how to install the HACMP/6000 software and its updates. It also covers the necessary customizations to the clhosts and clinfo.rc files.

- Chapter 6, “Cluster Environment Definition”

  The definition of the cluster, its nodes, and the network adapters for HACMP are given in this chapter. The example cluster is used for the definitions.

- Chapter 7, “Node Environment Definition”

  This chapter describes how to define application servers, resource groups, and resources belonging to those resource groups.

- Chapter 8, “Starting and Stopping Cluster Services”

  The options involved in starting and stopping the HACMP software on a machine are described here.

- Chapter 9, “Error Notification Tool”

Once the basic cluster has been set up and tested, error notification can be used to take special action upon the occurrence of specified errors in the AIX error log. The set of tools included in this book includes a tool that makes the setup and testing of these error notification methods quite easy.

- Chapter 10, “Event Customization Tool”

  This chapter describes a tool provided with the book that makes the customization of cluster events easier. It provides an example of using the tool.

- Chapter 11, “Cluster Documentation”

  The documentation tool provided with this book generates extensive documentation of a cluster node and cluster definitions. This documentation report can be used to allow a new administrator to understand the original setup of the cluster. This chapter describes how to run the documentation tool and generate a report.

- Appendix A, “Qualified Hardware for HACMP”

  This appendix includes the HAMATRIX document, which lists the tested and supported hardware for HACMP, as of the date of publication. This document is continually updated as new devices are introduced.

- Appendix B, “RS232 Serial Connection Cable”

  This appendix describes the options for buying or building the RS232 connection cable that is used to connect nodes with a non-TCP/IP network.

- Appendix C, “List of AIX Errors”

  This appendix provides a list of AIX errors that can be put into the AIX error log. It can be used as a reference in using the error notification tool.

- Appendix D, “Disk Setup in an HACMP Cluster”

  This appendix gives detailed descriptions of the cable requirements and other activities involved in connecting any of the supported shared disks for HACMP.

- Appendix E, “Example Cluster Planning Worksheets”

  This appendix includes completed cluster planning worksheets for the example cluster whose setup we describe in the document.

- Part 1, “Cluster Documentation Tool Report”

  This appendix includes a cluster documentation report, generated by the documentation tool included with this redbook.

Related Publications

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

- *HACMP/6000 Concepts and Facilities*, SC23-2699
- *HACMP/6000 Planning Guide*, SC23-2700
- *HACMP/6000 Installation Guide*, SC23-2701
- *HACMP/6000 Administration Guide*, SC23-2702
- *HACMP/6000 Troubleshooting Guide*, SC23-2703
- *HACMP/6000 Programming Locking Applications*, SC23-2704

- *HACMP/6000 Programming Client Applications*, SC23-2705
- *HACMP/6000 Master Index and Glossary*, SC23-2707
- *HACMP/6000 Licensed Program Specification*, GC23-2698
- *Common Diagnostics and Service Guide*, SA23-2687
- *RISC System/6000 System Overview and Planning*, GC23-2406

International Technical Support Organization Publications

- *HACMP/6000 Customization Examples*, SG24-4498
- *High Availability on the RISC System/6000 Family*, SG24-4551
- *A Practical Guide to the IBM 7135 RAID Array*, SG24-2565

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

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

To get a catalog of ITSO redbooks, VNET users may type:

```
TOOLS SENDTO WTSCPOK TOOLS REDBOOKS GET REDBOOKS CATALOG
```

A listing of all redbooks, sorted by category, may also be found on MKTTOOLS as ITSOCAT TXT. This package is updated monthly.

How to Order ITSO Redbooks

IBM employees in the USA may order ITSO books and CD-ROMs using PUBORDER. Customers in the USA may order by calling 1-800-879-2755 or by faxing 1-800-445-9269. Almost all major credit cards are accepted. Outside the USA, customers should contact their local IBM office. For guidance on ordering, send a PROFS note to BOOKSHOP at DKIBMVM1 or E-mail to bookshop@dk.ibm.com.

Customers may order hardcopy ITSO books individually or in customized sets, called BOFs, which relate to specific functions of interest. IBM employees and customers may also order ITSO books in online format on CD-ROM collections, which contain redbooks on a variety of products.

ITSO Redbooks on the World Wide Web (WWW)

Internet users may find information about redbooks on the ITSO World Wide Web home page. To access the ITSO Web pages, point your Web browser to the following URL:

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

IBM employees may access LIST3820s of redbooks as well. The internal Redbooks home page may be found at the following URL:

```
http://w3.itsc.pok.ibm.com/redbooks/redbooks.html
```

Acknowledgments

This project was designed and managed by:

David Thiessen
International Technical Support Organization, Austin Center

The authors of this document are:

Nadim Tabassum
IBM France

David Thiessen
International Technical Support Organization, Austin Center

The document is based on a version in the French language used in IBM France. The authors of the original document are:

C. Castagnier
IBM France

J. Redon
IBM France

Nadim Tabassum
IBM France

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:

Marcus Brewer
International Technical Support Organization, Austin Center

Chapter 1. Overview of the Tools

This document should be used in conjunction with the tools provided on the included diskette. To install the tools onto a system, use the following command:

```
# tar xvf /dev/rfd0
```

The tools are installed in the /usr/HACMP_ANSS directory.

All the tools are written to use this directory. If you wish to change this, it will involve a considerable effort on your part, and your scripts may not be in the same place in all sites where you use the tool.

The main subdirectories are:

tools This directory contains the tools provided to help you customize your environment. There is a subdirectory for each tool under this directory. Certain files which are common to all of the tools are also stored here.
DOC_TOOL - there are two tools here. The first, inventory, is used to obtain the state of the system before installing HACMP. This will also give you a list of any problems you may encounter due to different machines having similar logical volume names, SCSI ids, or other characteristics. The second tool, doc_dossier, produces a detailed description of your cluster configuration and should be run after installing HACMP. You can print out the report either in an ascii, VM or PostScript format.
ERROR_TOOL - this tool allows you to customize the handling of system errors.
EVENT_TOOL - this tool allows you to customize the actions taken in response to cluster events.

script This directory is not created at install time. It is created the first time one of the tools needs to write something into it. You should place all of your customized scripts here and this directory should never be deleted. Skeleton files are created here for certain events and errors; these should be tailored to suit your needs.

utils This directory contains site specific scripts which are created by the tools.

dessin This directory contains the files used to draw the cluster configuration.

backup This directory is created the first time it is called. It contains the output files for the tools when they are run.

Log files for the messages, errors and warnings generated by the customized scripts are stored in the directory /var/HACMP_ANSS/log. This directory is automatically created the first time that the tools are used. It contains two files which are created when they are first invoked. The files are called:

- hacmp.errlog
- hacmp.eventlog

As you use the tool, you will notice a French flavor in the variable names and file names. This has been preserved to recognize the heritage of the tools.

1.1 Installation Tips

Do not copy /usr/HACMP_ANSS from one machine in order to install the tools onto another machine. The script, utils and backup subdirectories will contain customized files which are specific for that machine.

To recover the tool for installation upon another machine, use the SAVE script in the /usr/HACMP_ANSS/tools directory, which has been specifically designed for this task, or use the original diskette if you still have it. To run this script (do not forget to insert a diskette) issue the command:

```
# /usr/HACMP_ANSS/tools/SAVE
```

Chapter 2. Inventory Tool

This tool examines the system configuration and determines if there are any points where we might have to pay particular attention. The shell script is called `inventory` and is found in the directory `/usr/HACMP_ANSS/tools/DOC_TOOL`.

The output file contains information on the configured adapters and disks. If you take this file to another system and run `inventory`, the tool will compare the output of the two files and indicate any potential points of conflict between the two systems.

2.1 Inventory - Communication adapters

This part of the `inventory` tool detects the presence of ethernet, token ring or FDDI adapters and gives the following:

- Slot number it is installed in
- Device name of the adapter

2.2 Inventory - Disks

This part of the `inventory` tool does the following:

- Lists the disk adapters
- Checks the SCSI ID of each adapter so you will know whether you you will have to change it (SCSI disks ONLY)
- Lists the disks connected to an adapter
- Lists the logical volumes (LVs) and indicates whether they are mirrored or not
- Checks that LV names and mount points are unique for each filesystem on the cluster nodes
- Checks that LV names are not trivial (like lv00 or lv01)

2.3 Output from the Inventory Tool

You will need a diskette and a printer, if you wish to have a hard copy of the output. The diskette is used to transfer the inventory produced on one node to another node. This allows the tool to identify any potential problems or conflicts between nodes.

If your machine does not have a floppy disk drive, then use `ftp` or `rcp` to transfer the files across to the other node.

If you do not have a printer connected to your machine, you can use the tool to save the output files on to a DOS or UNIX diskette. Then you can print the output from a PC or other UNIX or AIX machine.

All these options are presented by a menu after the inventory program has terminated.

2.4 Output Files

You can always examine the results which are presented on the screen. All output files are saved in the /tmp directory, with the name prefixed by HACMPmachine- and followed by hostname and a suffix indicating the type of output.

On a machine with the hostname jack, the files would be called:

```
HACMPmachine-jack-conf
HACMPmachine-jack-lv
HACMPmachine-jack-tty
```

2.5 Sample Configuration

Figure 1 on page 5 shows an inventory report generated by the inventory tool.

```
 6    6   6666    6666    66666  6    6    66    6    6  666666
 6    6  6    6  6         6    66   6   6  6   66  66  6       66666
 666666  6    6   6666     6    6 6  6  6    6  6 66 6  66666
 6    6  6    6       6    6    6  6 6  666666  6    6  6       66666
 6    6   6666    6666     6    6    6  6    6  6    6  666666

 6    6    66    66666     6    6    6
 66   6   6  6   6    6    6    66  66
 6 6  6  6    6  6    6    6    6 66 6
 6  6 6  666666  6    6    6    6    6
 6    6  6    6  66666     6    6    6

The following serial ports were found:
ADAP   ADDRESS
sa1  00-00-S1
sa2  00-00-S2

The following ttys are configured:
TTY       TERM        LOGIN       STOPS       BPC
tty0      ibm3151     enable       1          8
tty1      dumb        disable      1          8

The following network adapters were found:
ent0 00-00-0E

The scsi0 adapter has its SCSI ID set to id 7
and has the following disks connected to it:
ADAPT     DISK            ADDRESS    VOLUME GROUP
scsi0     hdisk0     00-00-0S-00     rootvg
scsi0     hdisk1     00-00-0S-40     nadvg
scsi0     hdisk2     00-00-0S-50     nadvg

Volume group rootvg contains the following logical volumes
VG NAME      LV NAME      TYPE     MOUNT POINT               MIRROR
rootvg       hd6        paging     N/A               no mirrored copies defined
rootvg       hd5        boot       /blv              no mirrored copies defined
rootvg       hd7        sysdump    /mnt              no mirrored copies defined
rootvg       hd8        jfslog     N/A               no mirrored copies defined
rootvg       hd4        jfs        /                 no mirrored copies defined
rootvg       hd2        jfs        /usr              no mirrored copies defined
rootvg       hd1        jfs        /home             no mirrored copies defined
rootvg       hd3        jfs        /tmp              no mirrored copies defined
rootvg       hd9var     jfs        /var              no mirrored copies defined
rootvg       lvtmp      jfs        /netview          no mirrored copies defined

Volume group nadvg contains the following logical volumes
VG NAME      LV NAME      TYPE     MOUNT POINT               MIRROR
nadvg        fslv00     jfs        /alpha            mirror 2 copies
nadvg        beta       jfs        /beta             mirror 2 copies
nadvg        gamma      jfs        /gamma            mirror 2 copies
nadvg        delta      jfs        /delta            mirror 2 copies
nadvg        nadlog     jfslog     N/A               mirror 2 copies
nadvg        zeta       jfs        N/A               mirror 3 copies
nadvg        theta      jfs        N/A               mirror 3 copies
nadvg        lv_netview jfs        /usr/OV           no mirrored copies defined
nadvg        lv_sm6000  jfs        /usr/adm/sm6000   no mirrored copies defined
```

Figure 1. Example of an inventory on a NODE

2.6 Example of Anomalies Report

An example of /tmp/HACMPmachine-anomalies is shown below. This file is produced as a result of running inventory on the second machine. You must already have copied across the results of running inventory on the first machine.

```
   66     66666   66666  666666  6    6   66666     6     6666   6    6
  6  6      6       6    6       66   6     6       6    6    6  66   6
 6    6     6       6    66666   6 6  6     6       6    6    6  6 6  6
 666666     6       6    6       6  6 6     6       6    6    6  6  6 6
 6    6     6       6    6       6   66     6       6    6    6  6   66
 6    6     6       6    666666  6    6     6       6     6666   6    6

ANOMALIES: CONFIGURATION INFORMATION
COMPARING THE TWO NODES

IDENTIFYING rs232 PORTS ON THE TWO NODES
NODE: jack - tty0 dumb       disable    1          8
NODE: nadim - tty1 dumb       disable    1          8

CHECKING THE SCSI ID's OF THE SHARED ADAPTERS
NODE: jack: The scsi0 adapter has its SCSI ID set to id 7
NODE: nadim: The scsi0 adapter has its SCSI ID set to id 7

CHECKING THE MOUNT POINTS
The /lll directory has the same mount point on the 2 nodes
The /mountp directory has the same mount point on the 2 nodes

CHECKING THE LOGICAL VOLUME NAMES
logical volume : zz has the same name on the 2 systems
logical volume lv00 has a non significant name on NODE: jack
```

Figure 2. Example of a /tmp/HACMPmachine-anomalies file

2.7 When to Run the Inventory Tool

The inventory can be run at any time. However, it is most useful to run it early in your setup process. Typically you would run the tool on each machine that will be a cluster node, before you have connected your shared disks and defined your shared volume groups.

Chapter 3. Setting up a Cluster

This chapter will begin to illustrate the setup of an HACMP cluster, using the set of tools provided with this document. This chapter, and those to follow, will cover:

- Planning Considerations
- Pre-Installation Activities
- Installing HACMP
- Cluster Environment Definition
- Node Environment Definition
- Starting and Stopping HACMP
- Error Notification Customization
- Event Customization
- Documenting your Cluster

Spread throughout our example will be descriptions of the correct times to run each of the various tools provided.

3.1 Cluster Description

We will now describe the cluster we are about to set up. This cluster will consist of two nodes, and will be set up in what is traditionally called a *Mutual Takeover* configuration. This is a configuration where each node serves a set of resources during normal operations, and each node provides backup for the other. There will also be a concurrent access volume group included. The cluster to be built is shown in Figure 3 on page 8.

Several observations should be made about this cluster:

- The cluster nodes are evenly matched 5XX model CPUs. This makes them good candidates for Mutual Takeover, since each node is able to handle an equal application load during normal operations.
- The main or *public network* is a Token-Ring network. Each node has two interfaces on this network, a service and a standby. Since we will be configuring each node to be able to take over the IP address of the other, each node will also have a *boot address* to be used on its service interface. i1.boot address This will allow the machine to boot and connect to the network without conflicts, when its *service address* has been taken over and is still active on the other node.
- There is a second network, an ethernet network called etnet1. This network will be defined to HACMP as a *private network*. As such, it will be used to carry Cluster Lock Manager traffic between nodes. A private network is highly recommended in any configuration using concurrent access. The private network has only service interfaces, and not standby interfaces. Standby interfaces can, of course, also be used in private networks, but since Cluster Lock Manager traffic automatically shifts to the public network if there is a private network failure, standby interfaces on a private network are not essential.

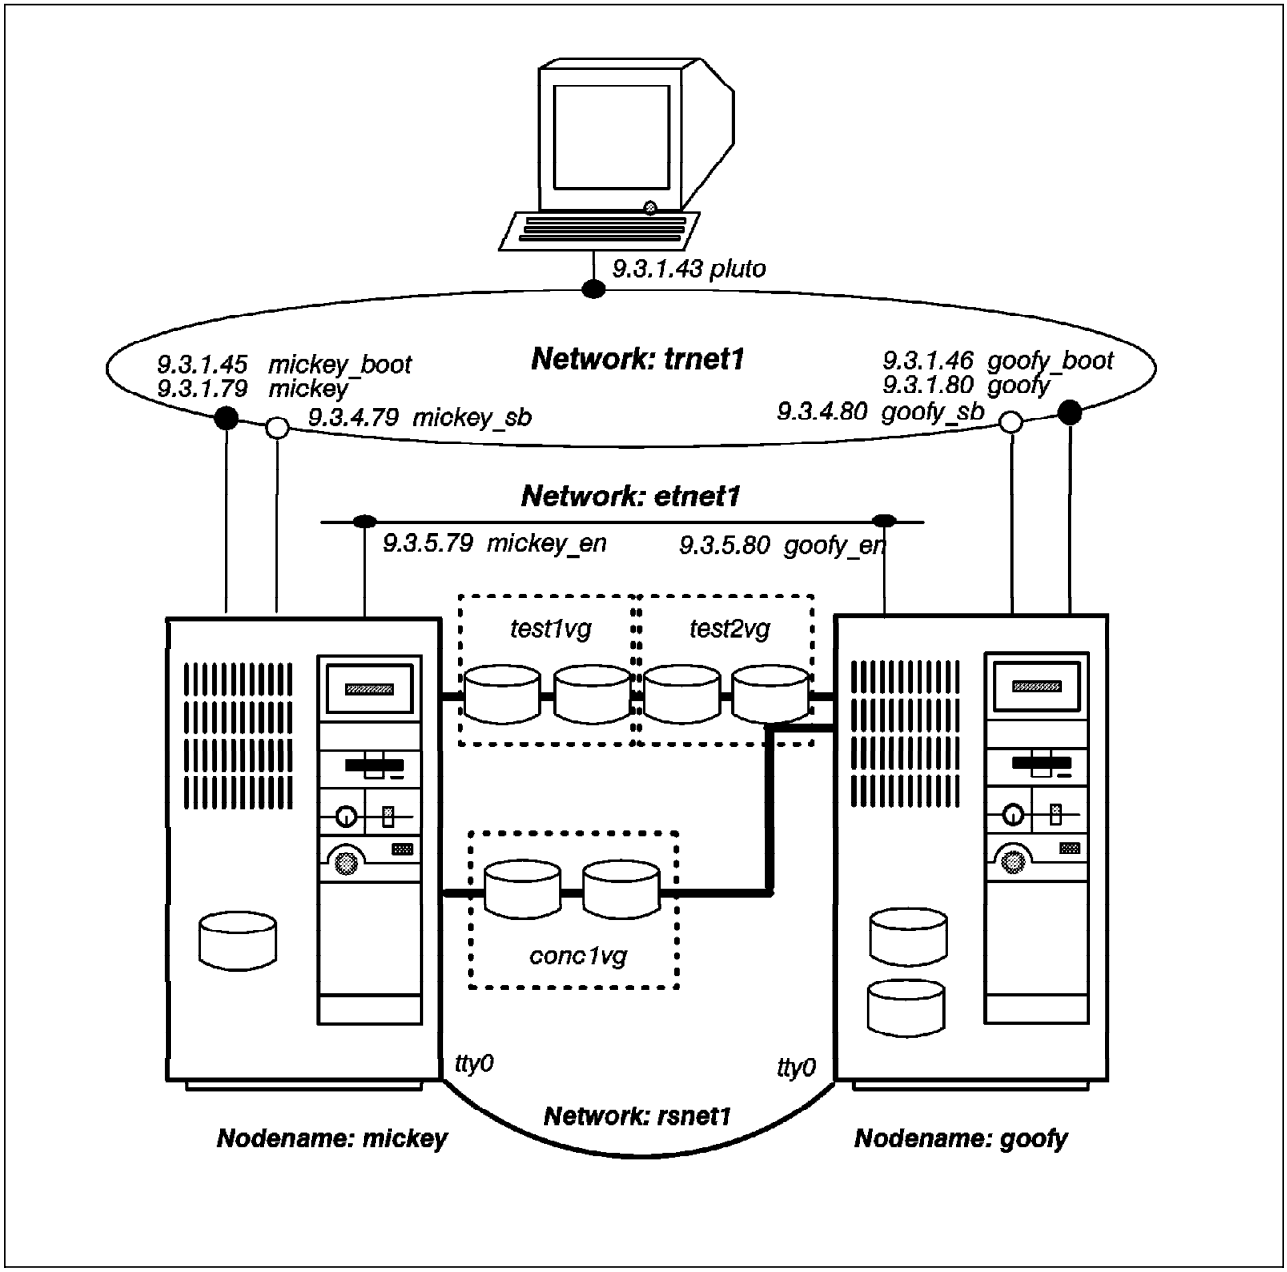

Figure 3. Cluster disney

- The cluster has IBM 9333 Serial disks as its shared disks. There are two 9333 subsystems connected. The first one includes four disk drives, which will be configured into two volume groups, each containing two disks. The second subsystem includes two disks, which will be contained in a single concurrent volume group. The node `mickey` has two 9333 disk adapters, each connected to one of the subsystems. The other node `goofy` has only one 9333 disk adapter, which is connected to both 9333 subsystems.
- There is also a raw RS232 link between native serial ports on the two nodes, who each have a tty device defined. This link will be defined as an HACMP network called `rsnet1`, and will be used so that the cluster can continue to send keepalive packets between nodes, even if the TCP/IP subsystems fail on one or more nodes.

- Node goofy has two internal disks in its rootvg volume group, while node mickey has only one. This will cause the shared disks to have different device names on each of the nodes. For example, one of the shared disks will be named hdisk1 on node mickey, and hdisk2 on node goofy. This is a common situation in clusters, and is nothing to worry about.
- There is a client system, connected on the token-ring network, called pluto. We will be installing the client component of the HACMP software on this system.

3.2 Planning Considerations

Depending on the type of hardware configuration you have in your cluster, you will have more or less planning considerations to deal with. If you are using SCSI disks as your shared disks, you will have more planning items to consider. Since we do not have shared SCSI disks in our example cluster, these concerns will not be ours in this setup, but we will deal with the planning items in this section. All cluster implementers must deal with planning items associated with:

- Networks
- Shared Disks
- Shared Volume Groups
- Planning Worksheets

3.2.1 Network Considerations

Every cluster should have one or more TCP/IP networks, and at least one non-TCP/IP network. The non-TCP/IP network allows keepalive packets to keep flowing from a node where the TCP/IP subsystem, but not the node itself, has failed. Either a raw RS232 link between systems, or a SCSI Target Mode connection can be used as a non-TCP/IP network. The setup of this network will be described later in this chapter.

3.2.1.1 TCP/IP Network Addresses

The following points must be considered when planning network addresses:

- The same subnet mask must be in use for all adapters on a node.
- Standby adapters must be on a different logical subnet from their service adapters.
- If a system will be having its service IP address taken over by another system, it must have a boot address configured. This boot address will be on the same logical subnet as the service address. The TCP/IP interface definition for the service adapter should be set to the boot address in this situation. If IP address takeover will not be used for this node, no boot address is necessary.

Please see the Planning Worksheets for our cluster in Appendix E, “Example Cluster Planning Worksheets” on page 131 to see how we have defined our adapters.

3.2.1.2 Hardware Address Takeover

HACMP can be configured to take over the hardware or MAC address of a network adapter, at the same time as it is taking over the IP address. If this facility is to be used, you must define, for each service interface that will have its address taken over, a dummy hardware address. This dummy address will be assumed by the adapter when it enters the cluster, and will be the hardware address that client systems associate with the system. This hardware address will then be moved, along with the IP address, whenever a failure in the cluster necessitates it.

This capability is only available for Token-Ring and ethernet networks. It allows you to have an IP address takeover, without having to refresh the ARP cache in each of the client systems. The relationship between IP address and MAC address remains constant throughout the takeover.

When you are defining a dummy hardware address, it is necessary for you to make sure that it does not conflict with any existing hardware address on the network. A good way to ensure this is to make your dummy address very close to the real hardware address of the adapter. For Token-Ring adapters, a convention for such an alternate hardware address is to change the first two digits of the real hardware address to 42. For ethernet adapters, there is no such convention. Many users will just change the last two digits of their adapter's address, and test with the `ping` command to make sure this address does not conflict.

3.2.2 Disk Adapter Considerations

The following considerations have to do with SCSI adapters only. If you are using 9333 Serial disks or 7133 SSA disks as your shared disks, you need not worry about any of these considerations. If you are using SCSI disks as your shared disks, you need to worry about several setup issues:

3.2.2.1 Termination

A SCSI bus must be terminated at each end. Normally, in a single system configuration, SCSI bus termination is done on the adapter at one end, by use of terminating resistor blocks. At the other end, the bus is terminated by a terminator plug, which is attached to the last device on the string.

In an HACMP cluster, you will have at least two and possibly more systems sharing the same set of SCSI disks. To be able to create a SCSI string, including both disk devices and SCSI adapters in systems, special Y-Cables are used. Also, the termination of the bus must be moved off the adapters themselves, and on to the Y-cables, to allow more than just two systems to share the bus.

Therefore, if you are using SCSI shared disks, you must use the correct Y-cables to connect them, and you must be sure to remove the terminating resistor blocks from each of your shared SCSI adapters. Depending on whether you are using 8-bit or 16-bit Fast/Wide adapters, the location of these terminating resistor blocks will be different.

There are pictures of the locations of these blocks on each of the adapters, as well as a full description of how to cable each of the types of shared disks with HACMP in Appendix D, "Disk Setup in an HACMP Cluster" on page 107.

3.2.2.2 SCSI IDs

It is mandatory, on a SCSI bus, that each device on the bus have a unique SCSI ID. Of course, everyone is used to making sure that each of the disk devices on a SCSI bus has a unique ID. In an HACMP cluster, you must also make sure that each of the adapters has a unique ID as well. Since SCSI adapters typically default to an ID of 7, this means you must change at least one.

It is highly recommended to change all SCSI adapter IDs to something other than 7. This is because certain recovery activities, including booting from diagnostic diskettes, return the SCSI adapters to ID 7, even though they might be configured for some other ID. If this is the case, an adapter under test could conflict with another adapter with that ID. Therefore, all shared SCSI adapter IDs should be changed from 7 to some other number. Since the highest ID always wins any arbitration for the SCSI bus, you should have all your adapters with the highest IDs on the bus.

There is a full description of how to change the SCSI ID on each of the supported types of SCSI adapters in Appendix D, “Disk Setup in an HACMP Cluster” on page 107.

3.2.2.3 Rebooting the Nodes

Whenever you have to reboot your cluster nodes, it is important that you do it one node at a time. If both nodes reach the point in their boot procedure where they are configuring the shared disks at the same time, you may have conflicts which will cause the disks not to be properly configured. This is why you should always first reboot one node, and wait until it has completed before rebooting the next node.

3.2.3 Shared Volume Group Considerations

There are several things to keep in mind when implementing shared volume groups. The special concerns have to do with naming and with major numbers.

3.2.3.1 Shared Volume Group Naming

Any shared volume group entity, including journaled filesystem logs (jfslogs), logical volumes, filesystems, and the volume groups themselves, must be explicitly named by you. If you allow the system to assign its default name for any of these items, you are most likely to have a naming conflict with an existing entity on one of the systems in the cluster.

Before you create any filesystems in your shared volume group, you should first create and explicitly name your jfslog. Once this is done, all filesystems you create in that volume group will use it.

Also, for any shared filesystems, you should not just create the filesystem, and allow the system to create the logical volume to contain it. This will allow the system to assign a logical volume name that is sure to conflict with something else in the cluster. Instead, first create the logical volume to contain the filesystem, giving it a unique name, and then create the filesystem on the logical volume. These procedures are shown later in our setup example.

3.2.3.2 Major Numbers

It is highly recommended to make sure that your shared volume groups have the same major number on each node. If you are exporting a shared filesystem through NFS to client systems, and a failure occurs, the client systems will only maintain their NFS mounts over the failure if the major number is the same. This is because NFS uses the major number as part of its file handle.

If you do not specify a major number when you create or import a shared volume group, the system will assign the next lowest unused number. Since different systems have different device configurations, it is almost certain that the next available number on each system will be different. You can check on the available major numbers on each system by running the `lvlstmajor` command. If you run this command on each node, and then choose a commonly available number to assign to your volume group, you will be OK.

A good recommendation is to use numbers much higher than any of the ones used in your system. For example, you might want to use numbers 60 and above to assign to your shared volume groups. We have found that, in upgrading to AIX Version 4.1, the system reserves many more major numbers than it did in AIX Version 3.2.5. If you use high numbers, you will not need to reassign your major numbers again if and when you upgrade to AIX Version 4.1.

3.2.4 Planning Worksheets

The *HACMP/6000 Planning Guide* includes a set of planning worksheets. These worksheets should be filled out when planning your cluster, before starting to set it up. These worksheets will force you to think through your planned configuration in detail, and make it much easier when it actually comes to doing the configuration. The completed worksheets for the cluster we will be setting up can be found in Appendix E, “Example Cluster Planning Worksheets” on page 131.

Chapter 4. Pre-Installation Activities

There are certain AIX configuration activities to be carried out before installing HACMP on your systems. These activities involve working on each of the systems that will become cluster nodes. They include preparing your network adapters, connecting your shared disks, and defining your shared volume groups.

4.1 Installing the Tools

Make sure that you have 2 MB free in the /usr filesystem. The tools will be installed into the directory /usr/HACMP_ANSS. The tools themselves take up less than 1 MB but they will create other directories and generate other programs.

Assuming you have the diskette included with this document, put it in your diskette drive, and issue the following commands:

```
# mkdir /usr/HACMP_ANSS
# tar -xvf/dev/fd0
```

If you do not have enough space in the /usr filesystem, and do not wish to make it bigger, you can make a separate filesystem for the tools by issuing the following commands:

```
# mklv -y'toolhacmp' rootvg 2
# crfs -v'jfs' -d'toolhacmp' -m'/usr/HACMP_ANSS' -A'yes' -p'rw' -t'no'
# mount /usr/HACMP_ANSS
# tar -xvf/dev/fd0
```

4.2 TCP/IP Configuration

The configuration of TCP/IP, before the installation of HACMP, involves:

- Configuration of adapters and hostnames
- Configuration of the /etc/hosts file
- Configuration of the /.rhosts file
- Testing

4.2.1 Adapter and Hostname Configuration

Now, each of the TCP/IP network adapters on your system must be defined to AIX. Use the worksheets you have prepared, or a diagram you have drawn of your cluster, like the one in Figure 3 on page 8, to refer to the network addresses you need.

Service and standby adapters should be configured. If you will be using a boot address, the service adapter should be configured to this address, rather than the service address.

It is recommended to configure the hostname of the system to be the same as the IP label for your service address, even if the IP address of the service adapter is initially set to the boot address.

You will issue the command smit mktcpip to take you to the panel where you will configure your service adapter:

```
                                Minimum Configuration & Startup

To Delete existing configuration data, please use Further Configuration menus

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

                                                                [Entry Fields]
* HOSTNAME                                                    [mickey]
* Internet ADDRESS (dotted decimal)                           [9.3.1.45]
  Network MASK (dotted decimal)                               [255.255.255.0]
* Network INTERFACE                                           tr0
  NAMESERVER
    Internet ADDRESS (dotted decimal)                         []
    DOMAIN Name                                               []
    Default GATEWAY Address                                   []
    (dotted decimal or symbolic name)
  RING Speed                                                  16                    +
  START Now                                                   yes                   +


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

Note that we have assigned a hostname of mickey, even though we have configured the IP address to be the boot address. If you are using a nameserver, be sure also to include the information about the server, and the domain, in this panel.

From here, we will use the command smit chinet to take us to the panel to configure the other network adapters. Here is the example for node mickey's standby adapter:

```
                    Change / Show a Token-Ring Network Interface

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

                                                              [Entry Fields]
  Network Interface Name                                     tr1
  INTERNET ADDRESS (dotted decimal)                         [9.3.4.79]
  Network MASK (hexadecimal or dotted decimal)              [255.255.255.0]
  Current STATE                                              up                  +
  Use Address Resolution Protocol (ARP)?                     yes                 +
  Enable Hardware LOOPBACK Mode?                             no                  +
  BROADCAST ADDRESS (dotted decimal)                        [9.3.4.255]
  Confine BROADCAST to LOCAL Token-Ring?                     no                  +



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

Continue with this for each of the TCP/IP network adapters on each of the nodes. If you have more than one network defined, also configure any service, boot, and standby adapters from those networks to TCP/IP.

4.2.2 Configuration of /etc/hosts File

Whether you are using nameserving or not, you will always want to include definitions for each of the cluster nodes' TCP/IP adapters in your /etc/hosts file. This will allow the cluster to continue working correctly even if your nameserver is lost.

You can either edit the /etc/hosts file directly, or use `smit hostent` to use SMIT for this purpose. Here is an example of the /etc/hosts definitions, configured for our example cluster:

```
# Cluster 1 - disney

9.3.1.45    mickey_boot
9.3.1.79    mickey
9.3.4.79    mickey_sb
9.3.5.79    mickey_en

9.3.1.46    goofy_boot
9.3.1.80    goofy
9.3.4.80    goofy_sb
9.3.5.80    goofy_en
```

Once you have created the /etc/hosts file on one system, you can use `ftp` to transfer it to each of your other cluster nodes.

4.2.3 Configuration of /.rhosts File

HACMP uses the /.rhosts file to allow it to carry out remote operations in other nodes. This is used for such things as synchronizing configurations between nodes, and running the cluster verification utility.

You should edit the /.rhosts file on the first node, and include each of the TCP/IP adapters on each of your cluster nodes. If you are using a nameserver, it is suggested to put each entry in its unqualified form, and also its fully qualified form, to allow the remote facilities to work correctly, whether the nameserver is available or not.

Here is an example of the /.rhosts file for our cluster:

```
mickey_boot
mickey
mickey_sb
mickey_en
goofy_boot
goofy
goofy_sb
goofy_en
mickey_boot.itsc.austin.ibm.com
mickey.itsc.austin.ibm.com
mickey_sb.itsc.austin.ibm.com
mickey_en.itsc.austin.ibm.com
goofy_boot.itsc.austin.ibm.com
goofy.itsc.austin.ibm.com
goofy_sb.itsc.austin.ibm.com
goofy_en.itsc.austin.ibm.com
```

Be sure the permissions on the /.rhosts file are set to 600; that is, read/write for root, and no access for anyone else. Again, once you have created this file correctly on one node, you can use ftp to transfer it to each of the others. Remember that any new files delivered by ftp will be set up with default permissions. You may need to sign on to each of the other nodes and change the permissions on the /.rhosts file.

4.2.4 Configuration of /etc/rc.net File

Unless you will be using your cluster node as a gateway or router, you should add the following statements to the end of the /etc/rc.net file:

```
/etc/no -o ipforwarding=0
/etc/no -o ipsendredirect=0
```

Again, if you are using your cluster nodes as gateways or routers, please skip this step.

4.2.5 Testing

Once you have completed this configuration, test it by using the ping command to contact each of your defined adapters, including standby adapters. If there is any problem here, do not continue until you have corrected it.

4.3 Non-TCP/IP Network Configuration

You will always want at least one non-TCP/IP network in your cluster. In our example, we will be using a raw RS232 link. If you are using SCSI differential shared disks, you have the option of using SCSI Target Mode communications as a network also. This will be described in this section also.

4.3.1 RS232 Link Configuration

The first set here is to connect the cable between serial ports on your systems. The cable can be bought from IBM or put together yourself, as described in Appendix B, “RS232 Serial Connection Cable” on page 97. Once you have connected the cable, you are ready for the next step.

4.3.1.1 Defining the tty Device

In most cases, you will use native serial ports on your systems for the RS232 link. This is what we are doing in our example, where we will be using the first native serial port, S1, on each node for our link.

Entering the command `smit mktty` will take you to the following panel:

```
                Add a TTY

  TTY type                                          tty
  TTY interface                                     rs232
  Description                                       Terminal asynchrone
  Parent adapter                                    sa0
* PORT number                                      [s1]
  BAUD rate                                        [9600]
  PARITY                                           [none]
  BITS per character                               [8]
  Number of STOP BITS                              [1]
  TERMINAL type                                    [dumb]
  STATE to be configured at boot time              [available]
  ...
  ...
  Enable LOGIN                                     disable
```

Use all the default settings, including leaving the `Enable LOGIN` field set to `disable`, and the `TERMINAL type` set to `dumb`. Take note of the tty device number returned by the SMIT panel, since you will need it later. If this is the first tty device defined, it will be /dev/tty0, which we will use in our example.

Do this definition on each of your nodes.

4.3.1.2 Testing the RS232 Link

Run the following command on the first node:

```
# stty < /dev/tty0
```

After you have entered the command, nothing should happen until you run the same command on the second node:

```
# stty < /dev/tty0
```

If the connection has been properly set up, you should now see the output of the stty command on both nodes.

Make sure that this is working correctly before proceeding.

4.3.2 SCSI Target Mode Configuration

We are not using shared SCSI differential disks in our example, and therefore will not be using SCSI target mode in our cluster, but a description of how to set it up is included here.

SCSI target mode connections can only be used with SCSI-2 Differential or Differential Fast/Wide adapters, and then only when the shared devices are not RAID arrays.

The inter-node communication (keepalive packets) used by HACMP to monitor the state of the cluster can also be carried out between SCSI adapters and can be used in place of (or along with) the RS232 serial network.

To enable the target mode capability, you need to modify the characteristics of the SCSI adapter. This can be done from the command line:

```
# chdev -l scsi2 -a tm='yes'
```

It can also be done through SMIT, by entering the command smit chgscsi. The following panel is presented:

```
                    Change/Show Characteristics of a SCSI Adapter


  SCSI adapter                                          scsi2
  Description                                           SCSI I/O Controller
  Status                                                Available
  Location                                              00-06
  Adapter card SCSI ID                                 [6]    +#
  BATTERY backed adapter                                no    +
  ...
  Enable TARGET MODE interface  ================>      yes   +
  Target Mode interface enabled                         no
  [PLUS...2]
```

A reboot is not necessary but you must rerun the configuration manager.

```
# smit device
    Configure Devices Added After IPL
```

Do the following command to find the name of the target mode SCSI link device:

```
# lsdev -Cc tmscsi
```

If this is the first link you have created, the device name will be tmscsi0. Note this name down, since it will be used in our testing and in HACMP configuration.

4.3.2.1 Testing a SCSI Target Mode Connection

Test the connection by carrying out the following steps. This example assumes that our target mode SCSI device created on each node is tmscsi0.

On the first node, enter the following command:

```
# cat < /dev/tmscsi0.tm
```

On the other node, enter the command:

```
# cat /etc/motd > /dev/tmscsi0.im
```

The contents of the /etc/motd file should be listed on the node where you entered the first command.

4.4 Connecting Shared Disks

Use the instructions included in Appendix D, “Disk Setup in an HACMP Cluster” on page 107 to connect your shared disks. There are instructions there for all kinds of shared disks supported by HACMP.

4.5 Defining Shared Volume Groups

Now you can create the shared volume groups and filesystems that will reside on the shared disk devices. Our configuration will have three volume groups. Volume group test1vg will be in a resource group owned by node mickey, volume group test2vg will be in another resource group owned by node goofy, and volume group conc1vg will be a concurrent volume group.

Each volume group contains two disks, and the logical volumes are mirrored from one to the other.

Creating the volume groups, logical volumes, and file systems shared by the nodes in an HACMP/6000 cluster requires that you perform steps on all nodes in the cluster. In general, you first define all the components on one node (in our example, this is node mickey) and then import the volume groups on the other nodes in the cluster (in our example, this is node goofy). This ensures that the ODM definitions of the shared components are the same on all nodes in the cluster.

Non-concurrent access environments typically use journaled file systems to manage data, while concurrent access environments use raw logical volumes.

Figure 4 on page 20 lists the steps you complete to define the shared LVM components for non-concurrent access environments.

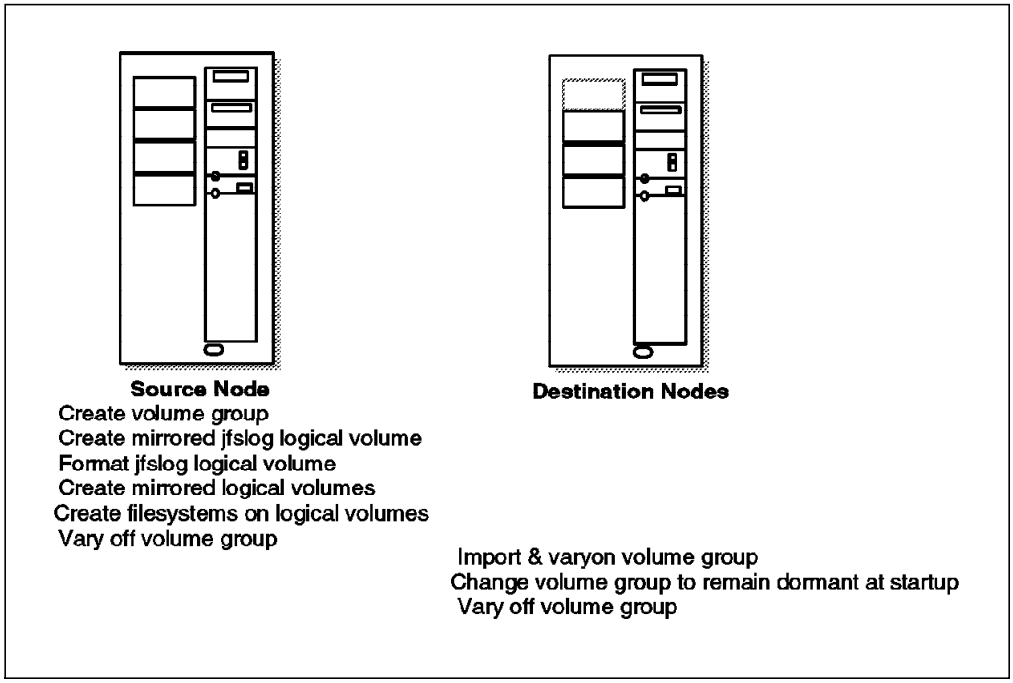

Figure 4. Defining Shared LVM Components for Non-Concurrent Access

For concurrent access, the steps are the same, if you omit those steps concerning the jfslog and filesystems.

4.5.1 Create Shared Volume Groups on First Node

Use the `smit mkvg` fastpath to create a shared volume group.

1. As root user on node `mickey` (the source node), enter `smit mkvg`:

```
                                Add a Volume Group

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

                                                      [Entry Fields]
  VOLUME GROUP name                                  [test1vg]
  Physical partition SIZE in megabytes                4                    +
* PHYSICAL VOLUME names                              [hdisk1 hdisk2]       +
  Activate volume group AUTOMATICALLY                 no                   +
    at system restart?
* ACTIVATE volume group after it is                  yes                  +
    created?
  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
```

Here, you provide the name of the new volume group, the disk devices to be included, and the major number to be assigned to it. It is also important to specify that you do not want the volume group activated (varied on) automatically at system restart, by changing the setting of that field to **no**.

The varyon of shared volume groups needs to be under the control of HACMP, so it is coordinated correctly.

Regardless of whether you intend to use NFS or not, it is good practice to specify a major number of the volume group. To do this, you must select a major number that is free on each node. Be sure to use the same major number on all nodes. Use the `lvlstmajor` command on each node to determine a free major number common to all nodes.

2. Because `test1vg` and `test2vg` contain mirrored disks, you can turn off quorum checking. On the command line, enter `smit chvg` and set quorum checking to **no**

```
                              Change a Volume Group

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

                                                      [Entry Fields]
* VOLUME GROUP name                                   test1vg
* 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
```

Now repeat the two steps above for volume group `test2vg`, using major number 61.

For our concurrent volume group `conc1vg`, with major number 62, repeat the two steps almost exactly, except that quorum protection must be left on for a concurrent volume group.

3. Varyon the three volume groups on node mickey:

```
# varyonvg test1vg
# varyonvg test2vg
# varyonvg conc1vg
```

4. Before you create any filesystems on the shared disk resources, you need to explicitly create the *jfslog logical volume*. This is so that you can give it a unique name of your own choosing, which is used on all nodes in the cluster to refer to the same log. If you do not do this, it is possible and likely that naming conflicts will arise between nodes in the cluster, depending on what user filesystems have already been created.

   Use SMIT to add the log logical volumes `loglvtest1` for the filesystems in volume group `test1vg`, and `loglvtest2` for the filesystems in volume group `test2vg`. Enter `smit mklv`, and select the volume group **test1vg** to which you are adding the first new jfslog logical volume.

```
                                Add a Logical Volume

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

[TOP]                                                   [Entry Fields]
  Logical volume NAME                                  [loglvtest1]
* VOLUME GROUP name                                    test1vg
* Number of LOGICAL PARTITIONS                        [1]                    #
  PHYSICAL VOLUME names                                [hdisk1 hdisk2]        +
  Logical volume TYPE                                  [jfslog]
  POSITION on physical volume                          midway                 +
  RANGE of physical volumes                            minimum                +
  MAXIMUM NUMBER of PHYSICAL VOLUMES                   []                     #
    to use for allocation
  Number of COPIES of each logical                     2                      +
    partition
  Mirror Write Consistency?                            yes                    +
  Allocate each logical partition copy                 yes                    +
    on a SEPARATE physical volume?
[MORE...9]

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

The fields that you need to change or add to are shown in **bold** type.

After you have created the jfslog logical volume, be sure to format the log logical volume with the following command:

```
# /usr/sbin/logform /dev/loglvtest1
logform: destroy /dev/loglvtest1 (y)?
```

Answer yes (y) to the prompt about whether to destroy the old version of the log.

Now create the log logical volume loglvtest2 for volume group test2vg and format the log, using the same procedure.

5. Now use SMIT to add the logical volumes lvtest1 in volume group test1vg and lvtest2 in volume group test2vg.

   It would be possible to create the filesystems directly, which would save some time. However, it is recommended to define the logical volume first, and then to add the filesystem on it. This procedure allows you set up mirroring and logical volume placement policy for performance. It also means you can give the logical volume a unique name.

   On node mickey, enter smit mklv, and select the volume group **test1vg**, to which you will be adding the new logical volume.

```
                                Add a Logical Volume

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

[TOP]                                                   [Entry Fields]
  Logical volume NAME                                  [lvtest1]
* VOLUME GROUP name                                     test1vg
* Number of LOGICAL PARTITIONS                        [20]                    #
  PHYSICAL VOLUME names                                [hdisk1 hdisk2]         +
  Logical volume TYPE                                  []
  POSITION on physical volume                           center                 +
  RANGE of physical volumes                             minimum                +
  MAXIMUM NUMBER of PHYSICAL VOLUMES                   []                      #
    to use for allocation
  Number of COPIES of each logical                      2                      +
    partition
  Mirror Write Consistency?                             yes                    +
  Allocate each logical partition copy                  yes                    +
    on a SEPARATE physical volume?
  RELOCATE the logical volume during                    yes                    +
    reorganization?
  Logical volume LABEL                                 []
  MAXIMUM NUMBER of LOGICAL PARTITIONS                 [128]
  Enable BAD BLOCK relocation?                          yes                    +
  SCHEDULING POLICY for writing logical                 sequential             +
    partition copies
  Enable WRITE VERIFY?                                  no                     +
  File containing ALLOCATION MAP                       []
[BOTTOM]

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

The **bold** type illustrates those fields that need to have data entered or modified. Notice that SCHEDULING POLICY has been set to **sequential**. This is the best policy to use for high availability, since it forces one mirrored write to complete before the other may start. In your own setup, you may elect to leave this option set to the default value of parallel to maximize disk write performance.

Again, repeat this procedure to create a 25 partition logical volume lvtest2 on volume group test2vg.

6. Now, create the filesystems on the logical volumes you have just defined. At the command line, you can enter the following fastpath: smit crjfslv. Our first filesystem is configured on the following panel:

```
       Add a Journaled File System on a Previously Defined Logical Volume


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

                                                        [Entry Fields]
* LOGICAL VOLUME name                                   lvtest1                +
* MOUNT POINT                                          [/test1]
  Mount AUTOMATICALLY at system restart?                no                     +
  PERMISSIONS                                           read/write             +
  Mount OPTIONS                                        []                      +
  Start Disk Accounting?                                no                     +




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

Repeat the above step to create the filesystem /test2 on logical volume lvtest2.

7. Mount the filesystems to check that creation has been successful.

```
# mount /test1
# mount /test2
```

8. If there are problems mounting the filesystems, there are two suggested actions to resolve them:

   a. Execute the fsck command on the filesystem.

   b. Edit the /etc/filesystems file, check the stanza for the filesystem, and make sure it is using the new jfslog you have created for that volume group. Also, make sure that the jfslog has been formatted correctly with the logform command.

   Assuming that the filesystems mounted without problems, now unmount them.

```
# umount /test1
# umount /test2
```

9. Now, create the logical volumes for our concurrent volume group conc1vg. From checking on the worksheet, you will see that we will be creating the following logical volumes:

   - conc1lv - 10 partitions - 2 copies
   - conc2lv - 7 partitions - 2 copies

10. Vary off the three volume groups.

```
# varyoffvg test1vg
# varyoffvg test2vg
# varyoffvg conc1vg
```

4.5.2 Import Shared Volume Groups to Second Node

The next step is to import the volume groups you have just created to node goofy. Login to node goofy as root and do the following steps:

1. Enter the fastpath command: smit importvg and fill out the fields as shown:

```
                              Import a Volume Group

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

                                                      [Entry Fields]
  VOLUME GROUP name                                  [test1vg]
* 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
```

2. Change the volume group to prevent automatic activation of test1vg at system restart and to turn off quorum checking. This must be done each time you import a volume group, since these options will reset to their defaults on each import. Enter smit chvg:

```
                              Change a Volume Group

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

                                                      [Entry Fields]
* VOLUME GROUP name                                  test1vg
* 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
```

3. Repeat the two steps above for volume group test2vg, using major number 61, and for conc1vg, using major number 62. For volume group conc1vg, leave quorum protection turned on, since this is a requirement for concurrent volume groups.

4. Vary on the volume groups and mount the filesystems on goofy to ensure that there are no problems.

Chapter 5. Installing the HACMP/6000 Software

The product is known as cluster on the AIX product tape. You can directly select the product using the / (find) option.

We may not install everything on each machine. Some machines may only require the client part or may not need the clvm.

5.1 On Cluster Nodes

On each node in the cluster, install the appropriate components of HACMP. From the panels you are led to from entering smit install, you will want to select the following:

```
>3.1.0.0 cluster
         cluster.client 03.01.00.00
         cluster.server 03.01.00.00
         cluster.clvm 03.01.00.00
```

Select your picks using F7.

In our example, we are selecting the option to install all components, including cluster.clvm which gives us the ability to do concurrent access.

If we were not running concurrent access, we would select cluster.server, which will automatically install cluster.client as a prerequisite.

5.2 On Cluster Clients

Here a client is considered to be a machine which is connected to the nodes through a network and accesses a highly available application on one of the cluster nodes. We restrict ourselves here to clients which are RISC System/6000s.

```
 3.1.0.0 cluster
  >       cluster.client 03.01.00.00
          cluster.server 03.01.00.00
          cluster.clvm 03.01.00.00
```

Select your picks using F7.

For non RS/6000 clients we can still carry out ARP cache refreshes using /usr/sbin/cluster/clinfo.rc.. Refer to Section 5.6, “Customizing the /usr/sbin/cluster/etc/clinfo.rc File” on page 29 to see how this is done.

5.3 Installing HACMP Updates

Now is the time to install the latest cumulative HACMP PTF fix from IBM. This should be done on both cluster nodes and client systems where you have installed the client portion of HACMP.

5.4 Loading the Concurrent Logical Volume Manager

Since we will be running with concurrent volume groups containing 9333 or SSA disks, we need to load the alternate Logical Volume Manager, called the Concurrent Logical Volume Manager (CLVM) which comes with HACMP. We will need to carry out this step on each node.

Loading the CLVM requires the following steps on each node:

1. Running the `cllvm -c concurrent` command
2. Running the command `bosboot -d /dev/ipldevice -a`
3. When the `bosboot` command completes, rebooting the system

Again, go through this procedure on each node.

Once the CLVM has been loaded as the active LVM, all continuing LVM administration can be done in the same way as with the standard LVM. The only exception is that the CLVM must be unloaded, and replaced with the standard IBM LVM before any AIX updates are applied to the system.

The procedure to reload the IBM standard LVM again is exactly the same as that shown above, except that the first step is to run the command `cllvm -c standard`. After the AIX updates have been loaded, the CLVM should be reloaded, using the above procedure, before returning the node to production in the cluster.

Again, these procedures are only required in an HACMP 3.1.1 cluster, if you have concurrent volume groups using 9333 or SSA disks. If you have concurrent volume groups using RAID arrays, you need not load the CLVM.

More information about loading the CLVM can be found in Chapter 6 of the *HACMP/6000 Installation Guide*.

5.5 Customizing the /usr/sbin/cluster/etc/clhosts File

On a client system, this file will be empty after the product installation. If you wish to use `clinfo`, then you must enter the boot and service addresses of each server node that this client should be able to contact.

On each server node, this file contains the loopback address which `clinfo` will use initially to acquire a cluster map. You should replace this with the boot and service addresses of all nodes in the cluster. On cluster nodes, this is not mandatory, but recommended.

Entries in this file can be one or the other of:

- symbolic names (IP labels)
- IP addresses

For example, you could add lines like :

```
mickey          # primary server
9.3.1.80        # backup server - goofy
```

5.6 Customizing the /usr/sbin/cluster/etc/clinfo.rc File

On each cluster node, if you have not implemented hardware address takeover, this file should contain a list of the IP addresses of its associated clients. This allows the node to ping the list of clients after a failure has occurred, so they can flush their ARP cache to reflect the new hardware address for a service adapter.

On each client system which uses the client portion of HACMP, this file should contain a list of the nodes with whom it communicates. Its default action is to flush the ARP cache, but you may want to extend this to execute your own programs. For example, you might want to display a window telling the user that the primary server is down and then display another message or window telling him that the backup server is now providing the services.

You will need to modify the following line in the file:

```
PING_CLIENT_LIST=""
```

These entries can be of the form:

- IP label (symbolic name)
- IP address

For instance:

```
PING_CLIENT_LIST="mickey goofy"
```

Clinfo is started automatically by the /etc/inittab file on cluster clients.

Chapter 6. Cluster Environment Definition

Defining the cluster environment involves making the following definitions:

- Cluster
- Cluster Nodes
- Network Adapters

These definitions can be entered from one node for the entire cluster. After this has been completed, the cluster environment definitions are synchronized from one node to all the others. Finally, the cluster environment should be verified, using the cluster verification utility, to ensure there are no errors before proceeding.

6.1 Defining the Cluster ID and Name

The first step is to create a cluster ID and name that uniquely identifies the cluster. This is necessary in case there is more than one cluster on a single physical network. Refer to your completed planning worksheets in Appendix E, "Example Cluster Planning Worksheets" on page 131 and complete the following steps to define the cluster ID and name.

1. Enter the `smit hacmp` command to display the system management menu for HACMP: The HACMP menu is the starting point for the definition and management of all HACMP characteristics and function.

```
                              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
```

2. Select **Manage Cluster Environment** and press Enter to display the following menu:

```
                    Manage Cluster Environment

Move cursor to desired item and press Enter.

  Configure Cluster
  Configure Nodes
  Configure Adapters
  Synchronize All Cluster Nodes
  Show Cluster Environment
  Configure Network Modules

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

3. Select **Configure Cluster** and press Enter to display the following menu:

```
                            Configure Cluster

Move cursor to desired item and press Enter.

  Add a Cluster Definition
  Change / Show Cluster Definition
  Remove Cluster Definition

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

4. Choose the **Add a Cluster Definition** option and press Enter to display the following panel.

```
                         Add a Cluster Definition

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

                                                    [Entry Fields]
  **NOTE: Cluster Manager MUST BE RESTARTED
          in order for changes to be acknowledged.**

* Cluster ID                                       [1]                    #
* Cluster Name                                     [disney]

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

5. Press Enter. The cluster ID and name are entered in HACMP's own configuration database managed by the ODM.

6. Press F3 to return to the Manage Cluster Environment screen. From here, we will move to the next stage, defining the cluster nodes.

6.2 Defining Nodes

Other parts of the cluster definition refer to the cluster nodes by their node names. In this section, we are simply defining the names that will identify each node in the cluster.

1. Select **Configure Nodes** on the Manage Cluster Environment screen to display the following menu:

```
                                Configure Nodes

Move cursor to desired item and press Enter.

  Add Cluster Nodes
  Change / Show Cluster Node Name
  Remove a Cluster Node




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

2. Choose the **Add Cluster Nodes** option and press Enter to display the following screen:

```
                             Add Cluster Nodes

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

                                                      [Entry Fields]
* Node Names                                         [mickey goofy]




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

   Remember to leave a space between names. If you use a duplicate name, an error message will be displayed. You need only to enter this information on one node, because you can later execute **Synchronize All Cluster Nodes** to propagate the information, using HACMP's Global ODM (GODM), to all other nodes configured in the cluster.

3. Press Enter to update HACMP's configuration database.

4. Press F3 to return to the Manage Cluster Environment screen. From here, we will move to the next stage, defining the network adapters to HACMP.

6.3 Defining Network Adapters

Having defined the node names, you can now proceed with defining the network adapters associated with each node. Again, you can define all the network adapters for all nodes on one node. You can later synchronize all the information to the other nodes' ODMs.

We shall use the values for our sample cluster. You should refer to the planning worksheets for TCP/IP and serial networks for your own cluster definitions. If you refer to Figure 3 on page 8, you will notice that both mickey and goofy contain two token-ring network adapters. One adapter is configured as a service adapter and the other is configured as a standby adapter. If the service adapter in one node fails, its standby adapter will be reconfigured by the Cluster Manager to take over that service adapter's IP address. If a node fails, the standby adapter in the surviving node will be reconfigured to take over the failed node's service IP address and masquerade as the failed node.

Notice also the RS232 connection between mickey and goofy. The RS232 link provides an additional path for keepalive (or heartbeat) packets and allows the Cluster Managers to continue communicating if the network fails. It is important to understand also that the RS232 network is not a TCP/IP network. Instead it uses HACMP's own protocol over the raw RS232 link.

Having this non-TCP/IP RS232 network is a very important requirement, since it provides us protection against two single points of failure:

1. The failure of the TCP/IP software subsystem
2. The failure of the single token-ring network

In either of these cases, if the RS232 network were not there, all keepalive traffic from node to node would stop, even though the nodes were still up and running. This is known as *node isolation*. If node isolation were to occur, mickey and goofy would both attempt to acquire their respective takeover resources. However, since the partner nodes would still be up and running, these attempts would fail, with the respective Cluster Managers endlessly attempting to reconfigure the cluster.

With the RS232 link in place, either of these failures would be interpreted as a network failure, instead of a node failure, allowing the administrator to take the appropriate action (restarting TCP/IP on a node, or fixing a network problem), without the cluster nodes trying to take over each other's resources inappropriately.

6.3.1 Defining mickey's Network Adapters

Complete the following steps to define mickey's network adapters:

1. Select **Configure Adapters** on the Manage Cluster Environments panel to display the following menu:

```
                              Configure Adapters

Move cursor to desired item and press Enter.

  Add an Adapter
  Change / Show an Adapter
  Remove an Adapter




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

2. Choose the **Add an Adapter** option. Press Enter to display the following panel, where you will fill out the fields for the service adapter:

```
                              Add an Adapter

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

                                                  [Entry Fields]
* Adapter Label                                   [mickey]
* Network Type                                    [token]                +
* Network Name                                    [trnet1]             +
* Network Attribute                                public                +
* Adapter Function                                 service               +
  Adapter Identifier                              [9.3.1.79]
  Adapter Hardware Address                        [0x42005aa8b484]
  Node Name                                       [mickey]               +




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

3. Press Enter to store the details in HACMP's configuration database.

   The following observations can be made about the fields to be filled in on this panel:

Adapter Label — This is the IP label of the adapter, which should be the same as the label you have defined in the /etc/hosts file and in your nameserver.

Network Type — If you list this field with F4, you will see the various Network Interface Modules (NIMs) available. There is a NIM for each type of network medium supported, as well as a Generic IP NIM. Since this adapter is on a token-ring network, we have selected the **token** NIM.

Network Name — This is an arbitrary name of your own choosing, to define to HACMP which of its adapters are on the same physical network. It is important that you use the same network name for all of the adapters on a physical network.

Network Attribute This field can either be set to public, private, or serial. A *public network* is one that is used by cluster nodes and client systems for access, as is this token-ring network. A *private network* is used for communications between cluster nodes only. The Cluster Lock Manager uses any private networks that are defined for its first choice to communicate between nodes. The most common reason to define a network as private is to reserve it for the exclusive use of the Cluster Lock Manager. A *serial network* is a non-TCP/IP network. This is the value you will define for your RS232 connection, and your SCSI Target Mode network if you have one.

Adapter Function This field can either be set to service, standby, or boot. A *service adapter* provides the IP address that is known to the users, and that is in use when the node is running HACMP and is part of the cluster. The *standby adapter*, as we have said before, is an adapter that is configured on a different subnet from the service adapter, and whose function is to be ready to take over the IP address of a failed service adapter in the same node, or the service adapter address of another failed node in the cluster. The *boot adapter* provides an alternate IP address to be used, instead of the service IP address, when the machine is booting up, and before HACMP Cluster Services are started. This address is used to avoid address conflicts in the network, because if the machine is booting after previously failing, its service IP address will already be in use, since it will have been taken over by the standby adapter on another node. A node rejoining the cluster will only be able to switch from its boot to its service address, after that service address has been released by the other node.

Adapter Identifier For a TCP/IP network adapter, this will be the IP address of the adapter. If you have already done your definitions in the `/etc/hosts` file, as you should have at this point, you do not have to fill in this field, and the system will find its value, based on the Adapter IP Label you have provided. For a non-TCP/IP (serial) network adapter, this will be the device name of the adapter, for instance `/dev/tty0` or `/dev/tmscsi0`.

Adapter Hardware Address This is an optional field. If you want HACMP to also move the hardware address of a service adapter to a standby adapter at the same time that it moves its IP address, you will want to fill in a hardware address here. This hardware address is of your own choosing, so you must make sure that it does not conflict with that of

any other adapter on your network. For token-ring adapters, the convention for an alternate hardware address is that the first two digits of the address are 42. In our example, we have found out the real hardware address of the adapter by issuing the command `lscfg -v -l tok0`. Our alternate hardware address is the same as the real address, except that we have changed the first two digits to 42. This ensures that there is not a conflict with any other adapter, since all real token-ring hardware address start with 10**...**. If you fill in an alternate hardware address here, HACMP will change the hardware address of the adapter from its real address which it has at boot time, to the alternate address, at the same time as it is changing the IP address from the boot address to the service address. If this is done, client users, who only know about the service address, will always have a constant relationship between the service IP address and its hardware address, even through adapter and node failures, and will have no need to flush their ARP caches when these failures occur. Alternate hardware address are only used with service adapters, since these are the only adapters that ever have their IP addresses taken over.

Node Name — This is the name of the node to which this adapter is connected. You can list the nodes that you have defined earlier with the F4 key, and choose the appropriate node.

4. Select the **Add an Adapter** option again. Press Enter to display the following panel and fill out the fields for the boot adapter:

```
                                Add an Adapter

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

                                                      [Entry Fields]
* Adapter Label                                      [mickey_boot]
* Network Type                                       [token]                 +
* Network Name                                       [trnet1]              +
* Network Attribute                                   public                 +
* Adapter Function                                   boot                    +
  Adapter Identifier                                 [9.3.1.45]
  Adapter Hardware Address                           []
  Node Name                                          [mickey]                 +

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

Notice that we have defined this adapter having the same network name as the service adapter. Also, you should note that the IP address for the boot adapter is on the same subnet as the service adapter. These two HACMP adapters, boot and service, actually represent different IP addresses to be used on the same physical adapter. In this case, token-ring adapter tok0 will start out on the boot IP address when the machine is first booted, and HACMP will switch the adapter's IP address to the service address (and the hardware address to the alternative address we have defined) when HACMP Cluster Services are started.

5. Press Enter to store the details in HACMP's configuration database.
6. Select the **Add an Adapter** option again. Press Enter and fill out the fields for the IP details for the standby adapter:

```
                                Add an Adapter

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

                                                      [Entry Fields]
* Adapter Label                                      [mickey_sb]
* Network Type                                       [token]                 +
* Network Name                                       [trnet1]              +
* Network Attribute                                   public                 +
* Adapter Function                                    standby                +
  Adapter Identifier                                 [9.3.4.79]
  Adapter Hardware Address                           []
  Node Name                                          [mickey]                +

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

Notice again that we have used the same network name, since this adapter is on the same physical network. We should also point out that this adapter has been configured on a different subnet from the boot and service adapter definitions. Our subnet mask was set earlier in the TCP/IP setup to 255.255.255.0.

7. Press Enter to store the details in HACMP's configuration database.
8. Select the **Add an Adapter** option again. Press Enter and fill out the details for the RS232 connection:

```
                                Add an Adapter

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

                                                      [Entry Fields]
* Adapter Label                                      [mickey_tty0]
* Network Type                                       [rs232]                +
* Network Name                                       [rsnet1]             +
* Network Attribute                                   serial                +
* Adapter Function                                    service               +
  Adapter Identifier                                 [/dev/tty0]
  Adapter Hardware Address                           []
  Node Name                                          [mickey]               +



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

Note here that we have chosen a different network type and network attribute, and assigned a different network name. Also, the adapter identifier is defined as the device name of the tty being used.

6.3.2 Defining goofy's Network Adapters

Repeat steps 2 on page 35 through 8 on page 38 to configure the adapters on goofy. Remember that all the configuration work can be done on one node because you can later synchronize this information to the other node(s) using HACMP's GODM facility.

Enter the service adapter details for goofy:

```
                                Add an Adapter

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

                                                      [Entry Fields]
* Adapter Label                                     [goofy]
* 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]                +



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

Here note that we have defined an alternate hardware address for this adapter also, which corresponds to the real hardware address of adapter tok0, with the first two digits changed to 42.

Enter the boot adapter details for goofy:

```
                                Add an Adapter

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

                                                      [Entry Fields]
* Adapter Label                                      [goofy_boot]
* Network Type                                       [token]                 +
* Network Name                                       [trnet1]              +
* Network Attribute                                   public                 +
* Adapter Function                                    boot                   +
  Adapter Identifier                                 [9.3.1.46]
  Adapter Hardware Address                           []
  Node Name                                          [goofy]                 +




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

Enter the IP details for goofy's standby adapter:

```
                                Add an Adapter

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

                                                      [Entry Fields]
* Adapter Label                                      [goofy_sb]
* Network Type                                       [token]                 +
* Network Name                                       [trnet1]              +
* Network Attribute                                   public                 +
* Adapter Function                                    standby                +
  Adapter Identifier                                 [9.3.4.80]
  Adapter Hardware Address                           []
  Node Name                                          [goofy]                 +




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

Enter the details for goofy's RS232 connection:

```
                                Add an Adapter

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

                                                      [Entry Fields]
* Adapter Label                                      [goofy_tty0]
* Network Type                                       [rs232]                 +
* Network Name                                       [rsnet1]              +
* Network Attribute                                   serial                 +
* Adapter Function                                    service                +
  Adapter Identifier                                 [/dev/tty0]
  Adapter Hardware Address                           []
  Node Name                                          [goofy]                 +




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

6.4 Synchronizing the Cluster Definition on All Nodes

The HACMP configuration database must be the same on each node in the cluster. If the definitions are not synchronized across the nodes, a run-time error message is generated at cluster startup time.

You will use the **Synchronize All Cluster Nodes** option on the Manage Cluster Environment panel to copy the cluster definition from mickey to goofy.

```
                            Manage Cluster Environment

Move cursor to desired item and press Enter.

  Configure Cluster
  Configure Nodes
  Configure Adapters
  Synchronize All Cluster Nodes
  Show Cluster Environment
  Configure Network Modules




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

1. Select the **Synchronize All Cluster Nodes** option on the Manage Cluster Environment menu and press Enter.

   SMIT responds: ARE YOU SURE?

2. Press Enter.

Note:

Before synchronizing the cluster definition, all nodes must be powered on, and the `/etc/hosts` and `/.rhosts` files must include all HACMP IP labels.

The cluster definition, including all node, adapter, and network module information, is copied from `mickey` to `goofy`.

For more information, refer to Chapter 8, Defining the Cluster Environment, in the *HACMP/6000 Installation Guide*.

Chapter 7. Node Environment Definition

This step entails telling HACMP how you would like it to behave when cluster events happen. Here you define the applications that will be managed by HACMP, and also the other resources, such as volume groups, filesystems, and IP addresses. By assigning node priorities, you also tell HACMP which node should take over the resources at what time.

The node environment definition stage involves three major steps:

- Defining application servers
- Defining resource groups and resources
- Verifying the cluster

7.1 Defining Application Servers

An *Application Server* defines a highly available application to HACMP. The definition consists of the following:

- Name
- Application start script
- Application stop script

Using this information, the application can be defined as a resource protected by HACMP. HACMP will then be able to start and stop the application at the appropriate time, and on the correct node.

Application Server start and stop scripts should be contained on the internal disks of each node, and must be kept in the same path location on each node. To define an Application Server, perform the following tasks:

1. At the command prompt, enter the SMIT fastpath `smit hacmp`. The following panel is presented:

```
                        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
```

2. Select **Manage Application Servers** to display the following screen:

```
                    Manage Application Servers

Move cursor to desired item and press Enter.

  Add an Application Server
  Change / Show an Application Server
  Remove an Application Server




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

3. Choose **Add an Application Server** to display the following screen:

```
                         Add an Application Server

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

                                                   [Entry Fields]
* Server Name                                     [mickeyapp1]
* Start Script                                    [/usr/local/mickey_start>
* Stop Script                                     [/usr/local/mickey_stop>




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

4. Enter an arbitrary Server Name, and then enter the full pathnames for the start and stop scripts. Remember that the start and stop scripts must reside on each participating cluster node. Our script names are:

   - /usr/local/mickey_start
   - /usr/local/mickey_stop

   Once this is done, an Application Server named mickeyapp1 has been defined, and can be included in a resource group to be controlled by HACMP.

You can now repeat a similar procedure to define an application server for goofy's application, called goofyapp1. Finally, you could create an application for the concurrent application, called concapp1.

7.2 Creating Resource Groups

In this section we shall go through the steps of defining two *cascading resource groups*, mickeyrg and goofyrg, and one *concurrent resource group*, concrg , to HACMP. Both nodes will participate in each resource group. Node mickey will have a higher priority for resource group mickeyrg and node goofy will have a higher priority for resource group goofyrg. In other words, mickey will own the resources in resource group mickeyrg, and will be backed up by goofy, while goofy will own the resources in resource group rg2, backed up by mickey. This is called *mutual takeover with cascading resources*.

Resource group mickeyrg will consist of the following resources:

- /test1 filesystem
- mickey's service IP address
- NFS export of the /test1 filesystem
- Application Server mickeyapp1

Resource group goofyrg will consist of the following resources:

- /test2 filesystem
- goofy's service IP address
- NFS export of the /test2 filesystem
- Application Server goofyapp1

As a final step, we will define our concurrent resource group concrg. Resource group concrg will consist of the following resources:

- logical volume conc1lv
- logical volume conc2lv
- Application Server concapp1

The steps required to set up this configuration of resource groups are as follows:

1. Configure the resource group mickeyrg on node mickey by using the SMIT fastpath command:

```
# smit cl_mng_res
```

   Then select **Add / Change / Show / Remove a Resource Group** from the following menu:

```
                              Manage Resource Groups

Move cursor to desired item and press Enter.

  Add / Change / Show / Remove a Resource Group
  Configure Resources for a Resource Group
  Configure Run Time Parameters




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

2. Select **Add a Resource Group** from the next menu:

```
                Add / Change / Show / Remove a Resource Group

Move cursor to desired item and press Enter.

  Add a Resource Group
  Change / Show a Resource Group
  Remove a Resource Group




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

3. In the panel that follows, fill out the fields as shown:

```
                            Add a Resource Group

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

                                                  [Entry Fields]
* Resource Group Name                            [mickeyrg]
* Node Relationship                               cascading              +
* Participating Node Names                       [mickey goofy]          +




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

   In the field Participating Node Names, be sure to name the highest priority node *first*. For resource group mickeyrg, this is mickey, since it is the owner. Other nodes participating then get named, in decreasing order of priority. In a two node cluster, there is only one other name, but in a larger cluster, you may have more than two nodes (but not necessarily all nodes) participating in any resource group.

4. Press Enter to store the information in HACMP's configuration database.

5. Press F3 twice to go back to the Manage Resource Groups panel. Select **Configure Resources for a Resource Group**.

```
                            Manage Resource Groups

Move cursor to desired item and press Enter.

  Add / Change / Show / Remove a Resource Group
  Configure Resources for a Resource Group
  Configure Run Time Parameters




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

6. The list that appears should show only one resource group, mickeyrg. Select this item.

```
                        Select a Resource Group

   Move cursor to desired item and press Enter.

     mickeyrg

   F1=Help                 F2=Refresh              F3=Cancel
   F8=Image                F10=Exit                Enter=Do
   /=Find                  n=Find Next
```

7. In the SMIT panel that follows, fill out the fields as shown. Make sure that the Inactive Takeover Activated and the 9333 Disk Fencing Activated fields are set to **false**.

```
                  Configure Resources for a Resource Group

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

                                                  [Entry Fields]
  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                            [mickeyapp1]                    +
  Miscellaneous Data                             []

  Inactive Takeover Activated                    false                    +
  9333 Disk Fencing Activated                    false                    +



F1=Help            F2=Refresh         F3=Cancel          F4=List
F5=Reset           F6=Command         F7=Edit            F8=Image
```

The following comments should be made about some of these parameters:

Service IP label — By filling in the label of **mickey** here, we are activating IP address takeover. If node mickey fails, its service IP address (and hardware address since we have defined it) will be transferred to the other node in the cluster. If we had left this field blank, there would be no IP address takeover from node mickey to node goofy.

Filesystems — Any filesystems that are filled in here will be mounted when a node takes over this resource group. The volume group that contains the filesystem will first be automatically varied on as well.

Filesystems to Export — Filesystems listed here will be NFS exported, so they can be mounted by NFS client systems or other nodes in the cluster.

Filesystems to NFS mount Filling in this field sets up what we call an *NFS cross mount*. Any filesystem defined in this field will be NFS mounted by all the participating nodes, other than the node that currently is holding the resource group. If the node holding the resource group fails, the next node to take over breaks its NFS mount of this filesystem, and mounts the filesystem itself as part of its takeover processing.

Volume Groups This field does not need to be filled out in our case, because HACMP will automatically discover which volume group it needs to vary on in order to mount the filesystem(s) we have defined. This field is there, so that we could specify one or more volume groups to vary on, in the case where there were no filesystems, but only raw logical volumes being used by our application.

Raw Disk PVIDs This field is very rarely used, but would be used in the case where an application is not using the logical volume manager at all, but is accessing its data directly from the hdisk devices. One example of this might be an application storing its data in a RAID-3 LUN. RAID-3 is not supported at all by the LVM, so an application using RAID-3 would have to read and write directly to the hdisk device.

Application Servers For any Application Servers that are defined here, HACMP will run their start scripts when a node takes over the resource group, and will run the stop script when that node leaves the cluster.

8. In the same way, set up the second resource group goofyrg.

```
# smit cl_mng_res
```

The following panel is displayed:

```
                          Manage Resource Groups

Move cursor to desired item and press Enter.

  Add / Change / Show / Remove a Resource Group
  Configure Resources for a Resource Group
  Configure Run Time Parameters

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

Select **Add / Change / Show / Remove a Resource Group**.

```
                Add / Change / Show / Remove a Resource Group

Move cursor to desired item and press Enter.

  Add a Resource Group
  Change / Show a Resource Group
  Remove a Resource Group




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

Select **Add a Resource Group**. On the resulting panel, fill in the fields, as shown below, to define your second resource group.

```
                              Add a Resource Group

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

                                                      [Entry Fields]
* Resource Group Name                                [goofyrg]
* Node Relationship                                   cascading              +
* Participating Node Names                           [goofy mickey]          +




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

Use F3 to go back to the Manage Resource Groups panel.

```
                            Manage Resource Groups

Move cursor to desired item and press Enter.

  Add / Change / Show / Remove a Resource Group
  Configure Resources for a Resource Group
  Configure Run Time Parameters




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

Select **Configure Resources for a Resource Group**.

```
                          Select a Resource Group

    Move cursor to desired item and press Enter.

      mickeyrg
      goofyrg

    F1=Help                 F2=Refresh              F3=Cancel
    F8=Image                F10=Exit                Enter=Do
    /=Find                  n=Find Next
```

Choose the resource group **goofyrg**.

```
                    Configure Resources for a Resource Group

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

                                                      [Entry Fields]
  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                                [goofyapp1]                    +
 Miscellaneous Data                                 []

 Inactive Takeover Activated                         false                  +
 9333 Disk Fencing Activated                         false                  +



F1=Help            F2=Refresh          F3=Cancel           F4=List
F5=Reset           F6=Command          F7=Edit             F8=Image
```

Fill in the appropriate fields, as shown above, and hit Enter to save the configuration.

9. Finally, we will set up our concurrent resource group concrg.

```
# smit cl_mng_res
```

The following panel is displayed:

```
                              Manage Resource Groups

Move cursor to desired item and press Enter.

  Add / Change / Show / Remove a Resource Group
  Configure Resources for a Resource Group
  Configure Run Time Parameters









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

Select **Add / Change / Show / Remove a Resource Group**.

```
                Add / Change / Show / Remove a Resource Group

Move cursor to desired item and press Enter.

  Add a Resource Group
  Change / Show a Resource Group
  Remove a Resource Group




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

Select **Add a Resource Group**. On the resulting panel, fill in the fields, as shown below, to define the concurrent resource group.

```
                              Add a Resource Group

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

                                                      [Entry Fields]
* Resource Group Name                                 [concrg]
* Node Relationship                                    concurrent              +
* Participating Node Names                            [mickey goofy]          +




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

Use F3 to go back to the Manage Resource Groups panel.

```
                            Manage Resource Groups

Move cursor to desired item and press Enter.

  Add / Change / Show / Remove a Resource Group
  Configure Resources for a Resource Group
  Configure Run Time Parameters




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

Select **Configure Resources for a Resource Group**.

```
                          Select a Resource Group

   Move cursor to desired item and press Enter.

     concrg
     goofyrg
     mickeyrg

   F1=Help                F2=Refresh             F3=Cancel
   F8=Image               F10=Exit               Enter=Do
   /=Find                 n=Find Next
```

Choose the resource group **concrg**.

```
                    Configure Resources for a Resource Group

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

                                                      [Entry Fields]
  Resource Group Name                                 concrg
  Node Relationship                                   concurrent
  Participating Node Names                            mickey goofy

 Service IP label                                    []                  +
 Filesystems                                         []              +
 Filesystems to Export                               []              +
 Filesystems to NFS mount                            []              +
 Volume Groups                                       []                  +
 Concurrent Volume groups                            [conc1vg]                   +
 Raw Disk PVIDs                                      []                      +
 Application Servers                                 [concapp1]                   +
 Miscellaneous Data                                  []

 Inactive Takeover Activated                          false                  +
 9333 Disk Fencing Activated                          false                  +



F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
```

Fill in the appropriate fields, as shown above, and hit Enter to save the configuration.

In a concurrent resource group, the only two resources to be defined are:

- Concurrent volume group - this gives access to the logical volumes
- Application server

10. The next job is to synchronize the node environment configuration to the other node. Hit F3 three times to return you to the Manage Node Environment panel, as shown below:

```
                                        Manage Node Environment

Move cursor to desired item and press Enter.

  Manage Resource Groups
  Change/Show Cluster Events
  Sync Node Environment




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

Select **Sync Node Environment**. You will see a series of messages, as the ODMs on the other node(s) are updated from the definitions on your node.

You can also synchronize the resource group configuration from the command line by executing the `/usr/sbin/cluster/diag/clconfig -s -r` command.

Note for HACMP Version 2.1 Users

For those users that have used HACMP Version 2.1, it is important for you to note that in HACMP/6000 Version 3.1 and HACMP 4.1 for AIX, the node environment must also be synchronized explicitly, along with the cluster environment. This is a change from HACMP Version 2.1, where the node environment was automatically synchronized by the Global ODM.

7.3 Verify Cluster Environment

Once you have completed the cluster and node environment definitions, you should verify that the node configurations are consistent and correct over the entire cluster. To verify the cluster enter the SMIT fastpath:

```
# smit hacmp
```

Select **Verify Environment** from the following panel:

```
                          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
```

The following panel is presented:

```
                        Verify Environment

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

                                              [Entry Fields]
    Verify Cluster Networks, 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
```

Take the default on this panel, which is to verify both the network configurations and the resource configurations. The Global ODM of HACMP will check the definitions on all nodes, to make sure they are correct and consistent. It will also check various AIX system parameters and system files, to make sure they are set correctly for HACMP, and will check any application server scripts you have defined, to make sure they are on all the nodes where they need to be, and that they are executable. You should see several verification messages, but the results should yield no errors. If you encounter errors, you must diagnose and rectify them before starting the cluster managers on each node. Failure to rectify verification errors will cause unpredictable results when the cluster starts.

Chapter 8. Starting and Stopping Cluster Services

Cluster nodes can be made to join and leave the cluster voluntarily by starting and stopping cluster services. There are various options available for both actions, controlling the immediate and future behavior of the node in the cluster.

8.1 Starting Cluster Services

Provided your verification has run without highlighting any errors, you are now ready to start cluster services on one node at a time. Each node should be able to finish its *node_up* processing, before another node is started.

To start cluster services on a node, issue the smit fastpath command `smit clstart`, to bring up the following panel:

```
                        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?                 false              +
  Startup Cluster Lock Services?                true               +
  Startup Cluster Information Daemon?           true               +




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

Here, you can select all the defaults, and hit Enter to start cluster services on the node. Since we are running a concurrent access environment in our example, we would want to change the last two fields to `true`.

Here are some comments on some of the fields:

Start now, on system restart or both

The recommended setting for this field is to **now**. If you set it to `system restart` or `both`, it will put a record into the `/etc/inittab` file, so that HACMP cluster services are started automatically on the machine each time it boots. This is not a very good idea, because it may result in a node trying to join the cluster before fixes have been fully tested, or at a time when the impact of resource group movement in the cluster is not desired.

It is much better to have explicit control over when cluster services are started on a node, and for that reason, the **now** setting is recommended.

Startup Cluster Lock Services?

Cluster Lock Services are, almost in all cases, only needed in a concurrent access configuration. The Cluster Lock Manager is normally used to control access to concurrently varied on volume groups. Therefore, we will want to change the setting to **true**, since we have a concurrent access configuration.

Startup Cluster Information Daemon?

The cluster information daemon, or `clinfo`, is the subsystem that manages the cluster information provided through the clinfo API to applications. This option would need to be set to `true` if you were going to be running applications directly on the cluster node that used the clinfo API. An example of such an application would be the cluster monitor `clstat`, which is provided as part of the product. If you are not running such an application, or are running such an application, but on a client machine, this option can be left with its default of `false`.

If you are running a clinfo application on a client machine, it gets its information from the *clsmuxpd* daemon on a cluster node, and does not need `clinfo` to be running on that cluster node.

When you start cluster services on a node, you will see a series of messages on the SMIT information panel, and then its status will switch to `OK`. This does not mean the cluster services startup is complete, however. To track the cluster processing, and to know when it is completed, you must watch the two main log files of HACMP:

- /var/adm/cluster.log

  This log file tracks the beginning and completion of each of the HACMP event scripts. Only when the `node_up_complete` event completes is the node finished its cluster processing.

- /tmp/hacmp.out

  This is a more detailed log file, as it logs each command of the HACMP event scripts as they are executing. In this case, you not only see the start and completion of each event, but also each command being executed in running those event scripts.

It is recommended to run the tail -f command against each of these log files when you start up nodes in the cluster, so that you can track the successful completion of events, and so that you can know when the processing is completed.

8.2 Stopping Cluster Services

To stop cluster services on a node, issue the smit fastpath command smit clstop, to bring up the following panel:

```
                          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
```

Here are some comments on the field choices:

Stop now, on system restart or both

If you select now, the default, HACMP will be stopped immediately, and no further action controlling future behavior will be taken. If you chose system restart or both, the system would also remove any automatic startup line for HACMP from the /etc/inittab file.

BROADCAST cluster shutdown?

Controls whether a broadcast message is sent to all users when HACMP is shut down on a node.

Shutdown mode

If you choose graceful, HACMP will be shut down on the machine, and any resources being held will be released. However, no other nodes in the cluster will take over the resources. This is a good option when you want to just shut down HACMP on all nodes, one at a time.

If you choose graceful with takeover, the HACMP software will be shut down and the resources released from the node. The next highest

priority node defined for the resource groups will then take over the appropriate resources.

If you choose forced, the HACMP software will be stopped on the node, but the resources that it is holding will be retained.

8.3 Testing the Cluster

It is highly recommended at this point, that you spend some time testing the operations of your cluster. You should try to test every conceivable failure, and make sure the cluster is reacting, and successfully dealing with them.

Chapter 9. Error Notification Tool

HACMP includes a menu-driven facility to customize the AIX error notification function. This allows you to run your own shell scripts in response to specified errors appearing in the AIX error log. To further ease the customization of the error notification object in the ODM (errnotify) which deals with both software and hardware errors, an error notification tool is provided on the diskette.

The shell script is called `error_select` and is found in the `/usr/HACMP_ANSS/tools/ERROR_TOOL` directory.

9.1 Description

Hardware and software errors, incidents and operator messages are logged in the AIX error log. To avoid the need for someone to periodically examine the error log in search of particular errors, we can configure *Error Notification Methods* to react automatically to the arrival of these errors. The errors that you will want to trap and treat will be dependent upon your installation.

The error notification tool will do the following:

- Create the templates for the scripts in the `script` subdirectory. These scripts can then be customized so that they react in the desired way to the arrival of errors. A possible example would be to promote a serial disk adapter failure to a node failure.
- Customize the relevant error notification objects in the ODM.
- Provide a test environment so that errors can be sent by you into the error log, without any real errors actually occurring. This will allow you to test your scripts. For example, we can generate SCSI_ERR3 without physically touching the SCSI adapter or the attached disks.

9.2 Error Notification Example

In our example cluster, we have two 9333 serial disk adapters on node `mickey`, but only one adapter on node `goofy`. Therefore, if the 9333 adapter on `goofy` fails, its users would be cut off from all the disks. However, since we have IP address takeover and disk takeover in our resource group definition, if we were to cause a node failure in this event, the users would be able to access the disks, still using the same IP address, through node `mickey`. Therefore, our error notification customization will send a warning message to the users, initiate a controlled HACMP shutdown with takeover, and then shutdown the machine itself.

Also, as well as sending mail to the root user on `goofy`, we want to send mail to our general system administrator, who is on another machine in the network.

The menu you will see when you run the error notification tool is shown below. The menu is preconfigured for those errors that have most often been customized in our experience. We have limited our choice to errors which are hardware and permanent, but you can add any AIX errors to this menu that you wish.

```
++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++
+                                                                    +
+               Choose one option at a time                          +
+          You can choose different errors successively              +
+                                                                    +
+               Enter: end (when you have finished)                  +
+                                                                    +
++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++
 1) end
 2) **************
 3) X25  - X25 adapter error
 4) DISK - SCSI disk error
 5) LVM  - LOGICAL VOLUME MANAGER error
 6) SCSI - SCSI adapter error
 7) TOK  - TOKEN RING adapter error
 8) EPOW - POWER SUPPLY problem
 9) FDDI - FDDI adapter error
10) SDA  - SERIAL disk ADAPTER error
11) SDC  - SERIAL disk CONTROLLER error
12) TMSCSI - SCSI network problem
Amongst this list, which errors would you like to treat:
```

We will make the following selection for our error:

```
Amongst this list, which error would you like to treat: 10
```

We could also choose more errors at the same time, if we wished. Here is what we will see on the screen:

```
**************************************
**   UPDATING ODM
**************************************
/usr/HACMP_ANSS/utils/error_SDA applied
*********************************************************
**   In order to delete your choice from the ODM         **
**   use error_del                                       **
*********************************************************
```

This procedure, as well as the procedures used to deselect the errors, (created automatically by the tool) are put into the utils subdirectory.

```
/usr/HACMP_ANSS/utils/error_SDA
```

The following routines, which will be executed as soon as the relevant error is logged in the error log, will be automatically created in the /usr/HACMP_ANSS/script subdirectory.

```
error_SDA
error_NOTIFICATION
```

It is up to you to modify these scripts so that they behave as you require. As they are created by the tool, they are just empty template scripts.

The error_NOTIFICATION script, which is automatically invoked by the error_SDA script, logs the incident in the /var/HACMP_ANSS/log/hacmp.errlog file and sends a mail message to the root user.

Here is a listing of the error_SDA script, as we have modified it to our requirements:

```
#!/bin/ksh
################################################################################
# Written by: AUTOMATE
# Last modification by *** who ***
#
#  script:  error_SDA
#  parameters: 8 parameters (documented in error_NOTIFICATION)
#
# ARGUMENTS received :
# "sequence number in the error log = $1"
# "error ID                         = $2"
# "error class                      = $3"
# "error type                       = $4"
# "alert flag                       = $5"
# "resource name                    = $6"
# "resource type                    = $7"
# "resource class                   = $8"
# "error label                      = $9"
################################################################################
# Variables:
 . /usr/HACMP_ANSS/tools/tool_var
 STATUS=0
 ( echo "n=error_SDA==============date"
 echo "ERROR DETECTED: error_SDA") |tee -a $ERREURS/hacmp.errlog> /dev/console
  . $SCRIPTS/error_NOTIFICATION
###################### START OF CUSTOMIZATION ##############################
#
LOCALNODENAME=$(/usr/sbin/cluster/utilities/get_local_nodename)
mail -s "Error Alert" sysadm@theboss.company.com << END
An error has been detected on the HACMP cluster node $LOCALNODENAME
look at the $LOG file on the node.
DEVICE = $6
ADAPTER = $8

The system will be shut down and the users moved to a backup node.
END


wall "System will be shutting Down in 20 Seconds. Please log off now.
You will be able to login to your application again within 5 minutes."
sleep 20

#  This command does a shutdown with takeover of HACMP

/usr/sbin/cluster/utilities/clstop -y '-N'  '-gr'
sleep 5

#  We now want to shutdown the machine, until our administrator can
#  investigate the problem.

/etc/shutdown -Fr

######################  END OF CUSTOMIZATION  ##############################
return $STATUS
```

The error_NOTIFICATION script, automatically created along with error_SDA in the script subdirectory, looks like this:

```
#!/bin/ksh
######################################################################
#
#   name             : error_NOTIFICATION
#   INPUT paremeters : $1 to $8 sent by errpt
#   Description      : called by each error, sends a message
#                                  into hacmp.errlog
######################################################################
# Variables:
 . /usr/HACMP_ANSS/tools/tool_var
STATUS=0
G=$(tput smso)
F=$(tput rmso)
LOG="$ERREURS/hacmp.errlog"
#############################################################
#   main
#############################################################
(print "************ Source and cause of error ****************"
print "HOSTNAME=$(hostname)  DATE=$(date)"
print "sequence number in error log = $1"
print "error ID                     = $2"
print "error class                  = $3"
print "error type                   = $4"
print "alert flag                   = $5"
print "resource name                = $6"
print "resource type                = $7"
print "resource class               = $8"
print "error label                  = $9") >> $LOG
######################################################################
# DO NOT FORGET TO set TO_WHOM in error_MAIL
 . /usr/HACMP_ANSS/tools/ERROR_TOOL/error_MAIL $1 $2 $3 $4 $5 $6 $7 $8 $9
######################################################################
# DO NOT FORGET TO set QUEUE in error_PRINT
# . /usr/HACMP_ANSS/tools/ERROR_TOOL/error_PRINT $1 $2 $3 $4 $5 $6 $7 $8 $9
######################################################################

return $STATUS
```

The only customization required to this script might be to uncomment the line near the end that will cause a record of the error to be printed to the printer of your choice.

The /usr/HACMP_ANSS/tools/ERROR_TOOL/error_MAIL script, in its default form, will send mail to the root user on the system on which the error occurs. This could also be changed as required. The script is shown below:

```
#!/bin/ksh
# this script is executed if it has been uncommented in
# error_NOTIFICATION
#
#     variable: TO_WHOM should be set to the name of a user
#                and should be in the form
#               "user" or "user@hostname"
######################################################################
. /usr/HACMP_ANSS/tools/tool_var
TO_WHOM="root"
LOCALNODENAME=$(/usr/sbin/cluster/utilities/get_local_nodename)
mail $TO_WHOM << END
An error has been detected on the HACMP cluster node $LOCALNODENAME
look at the $LOG file
DEVICE = $6
ADAPTER = $8
END
```

Finally, if you wish to use the printing option, you will need to set the QUEUE variable in the /usr/HACMP_ANSS/tools/ERROR_TOOL/error_PRINT script to the name of a valid print queue for your system. The script is shown below:

```
#!/bin/ksh
# this script is executed if it has been uncommented in
# error_NOTIFICATION
#
#     variable: QUEUE should be set to a local or remote print queue
#               which has been defined in /etc/qconfig
######################################################################
QUEUE="NONE"
if [ $QUEUE = 'NONE' ]
then
    FILE_CIBLEE=''
else
    FILE_CIBLEE="-P $QUEUE"
fi
(banner 'Machine:' $(hostname )
print "=============================================================================="
print "$(date)"
print "=============================================================================="
print "refer to $LOG and look at errpt"
banner "error" "on" "device" "$6" )  | qprt  $FILE_CIBLEE
######################################################################
```

9.2.1 Checking the ODM

We will just do a check of the ODM to make sure that the error notification method has been set up correctly. Issue the SMIT fastpath command smit hacmp, and select the following options in the SMIT panels:

```
Cluster RAS Support
    Error Notification
          Change/Show a Notify Method
```

Our error notification tool actually set up two error notification methods, for the errors sda_err1 and sda_err3. If we choose the first one, the following panel is presented:

```
                          Change/Show a Notify Method

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

                                                      [Entry Fields]
* Notification Object Name                            sda_err1
* Persistence across system restart?                  Yes                    +
  Process ID for use by Notify Method                 [0]                    +#
  Select Error Class                                   All                    +
  Select Error Type                                    All                    +
  Match ALERTable errors?                              All                    +
  Select Error Label                                  [SDA_ERR1]              +
  Resource Name                                       []                      +
  Resource Class                                      []                      +
  Resource Type                                       []                      +
* Notify Method                                       [/usr/HACMP_ANSS/script>


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

Once we have customized these scripts as we want them, and have checked that they are correctly in the ODM, we are able to test the error notification method, simulating the actual error with the error testing tool.

9.3 Testing the Error Scripts

We can test the error handling scripts that we have created by running the /usr/HACMP_ANSS/tools/error_test script. This will send the required error to the AIX error log.

The menu that you will see when you start up error_test is shown below. As well as testing your scripts, this menu can be used during the acceptance testing phase to generate errors, without having to try to simulate them by pulling adapters and cables.

```
++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++
+                                                                    +
+               MENU: Testing errors                                 +
+                                                                    +
+                Choose one option at a time                         +
+          You can choose different errors successively              +
+                                                                    +
+                Enter: end (when you have finished)                 +
+                                                                    +
++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++
1) end
2) SDA_ERR1
3) SDA_ERR3
Which of the above errors would you like to generate:
```

If you wanted to run error_test to simulate SDA_ERR1, then you would do the following:

```
Which of the above errors would you like to generate: 2
```

You will have to enter the adapter for which you wish to simulate the error.

```
For which device are you simulating this error
For example enter: scsi2 hdisk4 ent0
The defective device is: serdasda0
```

Here is an example of what you will see on your screen:

```
The defective unit is: serdasda0

  Error id : b135ae8b
B135AE8B 1214112795 P H  serdasda0      STORAGE SUBSYSTEM FAILURE
FEC3157O 1213144095 P H  serdasda0      UNDETERMINED ERROR
B135AE8B 1213141195 P H  serdasda0      STORAGE SUBSYSTEM FAILURE
B135AE8B 1213120895 P H  serdasda0      STORAGE SUBSYSTEM FAILURE
FEC3157O 1213115495 P H  serdasda0      UNDETERMINED ERROR
B135AE8B 1213114095 P H  serdasda0      STORAGE SUBSYSTEM FAILURE
FEC3157O 1213104695 P H  serdasda0      UNDETERMINED ERROR
B135AE8B 1213101995 P H  serdasda0      STORAGE SUBSYSTEM FAILURE
FEC3157O 1212180795 P H  serdasda0      UNDETERMINED ERROR
B135AE8B 1212180595 P H  serdasda0      STORAGE SUBSYSTEM FAILURE
B135AE8B 1212175595 P H  serdasda0      STORAGE SUBSYSTEM FAILURE
BAECC981 1128181495 P H  serdasda0      MICROCODE PROGRAM ERROR
```

Each time this error is generated, the following entry will be added to the /var/HACMP_ANSS/log/hacmp.errlog file. This file should be checked periodically, since it will grow over time. The entry added is formatted by the error_NOTIFICATION program which can also send mail messages if desired.

```
=error_SDA===============Wed Dec 13 11:40:55 CST 1995
ERROR DETECTED: error_SDA
************ Source and cause of error ****************
HOSTNAME=goofy  DATE=Wed Dec 13 11:40:55 CST 1995
sequence number in error log = 1790
error ID                     = 0xb135ae8b
error class                  = H
error type                   = PERM
alert flag                   = TRUE
resource name                = serdasda0
resource type                = serdasda
resource class               = adapter
error label                  = SDA_ERR1
```

At the same time as the hacmp.errlog is being updated, the error_SDA shell script will be executed, carrying out whatever instructions you have added there.

For more information about error notification refer to the *AIX Problem Solving Guide*.

9.4 Deleting Error Notification Routines

You may decide that you no longer wish to take special action for a particular error. The procedures necessary to do this have been provided as part of the tool.

The script to use is called error_del. On running this script, the following menu will appear on the screen:

```
++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++
+                                                                    +
+          REMOVING AN ERROR NOTIFICATION OBJECT CLASS               +
+                                                                    +
+                Choose one option at a time                         +
+          You can remove different errors successively              +
+                                                                    +
+                Enter: end (when you have finished)                 +
+                                                                    +
++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++
1) end
2) SDA
Amongst this list, which errors would you like to remove: 2
```

Suppose you choose number **2**. The errnotify object class within ODM will automatically be modified, deleting the entry for the treatment of errors generated by the failure of the 9333 serial disk adapter.

The error_SDA script will be removed from the script subdirectory. The script is not actually deleted. Rather, it is moved to the backup subdirectory and its name is suffixed with YYYYMMDDhhmmss.

Chapter 10. Event Customization Tool

This tool helps in the customization of HACMP events. The main script is called `event_select` and is found in the `/usr/HACMP_ANSS/tools/EVENT_TOOL` directory.

10.1 Description

HACMP constantly surveys the states of the nodes in the cluster and at any given moment knows if:

- A node has failed
- A node has come up and has rejoined the cluster

Sometimes you need to customize HACMP′s reactions to an event because the event script, as provided with HACMP, does not fulfill your needs. For instance, you may have some of the following requirements:

- A node goes down. The cluster clients access this node through X.25. What must I do on the backup machine so that HACMP will correctly restart all the applications?
- A node goes down. The database has also crashed. What procedures do I have to run (rollback, redologs) before restarting the application on the backup machine?
- A node goes down. How do I recover the print jobs and cron jobs?

HACMP handles all changes to the cluster with cluster events. There are two types of events:

- Primary Events - 14 of them, called by the cluster manager
- Secondary or Sub Events - 16 of them, called by primary event scripts

A short description of each of the events is given below.

10.2 Primary Events

Event	Cause and action
config_too_long	Sends a periodic console message when a node has been in reconfiguration for more than six minutes.
fail_standby	Sends a console message when a standby adapter fails or is no longer available because it has been used to take over the IP address of another adapter.
join_standby	Sends a console message when a standby adapter becomes available.
network_down	Occurs when the cluster determines that a network has failed. The event script provided takes no default action, since the appropriate action will be site/LAN specific.
network_down_complete	Occurs only after a `network_down` event has successfully completed. The event script provided takes no default action, since the appropriate action will be site/LAN specific.

network_up	Occurs when the cluster determines that a network has become available. The event script provided takes no default action, since the appropriate action will be site/LAN specific.
network_up_complete	Occurs only after a network_up event has successfully completed. The event script provided takes no default action, since the action will be site/LAN specific.
node_down	Occurs when a node is detaching from the cluster, either voluntarily or due to a failure. Depending on whether the node is local or remote, either the node_down_local or node_down_remote sub event is called.
node_down_complete	Occurs only after a node_down event has successfully completed. Depending on whether the node is local or remote, either the node_down_local_complete or node_down_remote_complete sub event is called.
node_up	Occurs when a node is joining the cluster. Depending on whether the node is local or remote, either the node_up_local or node_up_remote sub event is called.
node_up_complete	Occurs only after a node_up event has successfully completed. Depending on whether the node is local or remote, either the node_up_local_complete or node_up_remote_complete sub event is called.
swap_adapter	Exchanges or swaps the IP addresses of two network interfaces. NIS and name serving are temporarily turned off during this event.
swap_adapter_complete	Occurs only after a swap_adapter event has successfully completed. Ensures that the local ARP cache is updated by deleting entries and pinging cluster IP addresses.
event_error	Occurs when an HACMP event script fails for some reason.

10.3 Secondary or Sub Events

Event	Cause and action
acquire_service_addr	Configures boot address to the corresponding service address and starts TCP/IP servers and network daemons by running the telinit -a command. HACMP modifies the /etc/inittab file by setting all the TCP/IP related startup records to a run level of a.
acquire_takeover_addr	Acquires takeover IP address by checking configured standby addresses and swapping them with failed service addresses.
get_disk_vg_fs	Acquire disk, volume group and file system resources as part of takeover.

node_down_local Releases resources taken from a remote node, stops application servers, releases a service address taken from a remote node, releases concurrent volume groups, unmounts file systems and reconfigures the node to its boot address.

node_down_local_complete Instructs the cluster manager to exit when the local node has completed detaching from the cluster. This event only occurs after a `node_down_local` event has successfully completed.

node_down_remote Unmounts any NFS file systems and places a concurrent volume group in non-concurrent mode when the local node is the only surviving node in the cluster. If the failed node did not go down gracefully, acquires a failed nodes resources: file systems, volume groups and disks and service address.

node_down_remote_complete Starts takeover application servers if the remote node did not go down gracefully. This event only occurs after a `node_down_remote` event has successfully completed.

node_up_local When the local node attaches to the cluster: acquires the service address, clears the application server file, acquires file systems, volume groups and disks resources, exports file systems and either activates concurrent volume groups or puts them into concurrent mode depending upon the status of the remote node(s).

node_up_local_complete Starts application servers and then checks to see if an inactive takeover is needed. This event only occurs after a `node_up_local` event has successfully completed.

node_up_remote Causes the local node to release all resources taken from the remote node and to place the concurrent volume groups into concurrent mode.

node_up_remote_complete Allows the local node to do an NFS mount only after the remote node is completely up. This event only occurs after a `node_up_remote` event has successfully completed.

release_service_addr Detaches the service address and reconfigures to its boot address.

release_takeover_addr Identifies a takeover address to be released because a standby adapter on the local node is masquerading as the service address of the remote node. Reconfigures the local standby into its original role.

release_vg_fs Releases volume groups and file systems that the local node took from the remote node.

start_server Starts application servers.

stop_server Stops application servers.

10.4 How the Event Customization Tool Works

Each of the HACMP events has a corresponding shell script in the /usr/sbin/cluster/events directory. Some of these shell scripts have no default action defined but are given as frameworks for you to fill in and customize as you wish.

When the cluster manager detects an event, it will run the associated script. This script is defined within the ODM by the HACMPevent object class found in /etc/objrepos/HACMPevent. The ODM entries for the first 3 events (before any modifications) are shown below:

```
HACMPevent:
    name = "swap_adapter"
    desc = "Swap adapter event happens. Swapping adapter."
    setno = 0
    msgno = 0
    catalog = ""
    cmd = "/usr/sbin/cluster/samples/swap_adapter"
    notify = ""
    pre = ""
    post = ""
    recv = ""
    count = 0

HACMPevent:
    name = "swap_adapter_complete"
    desc = "Swap adapter event completed."
    setno = 0
    msgno = 0
    catalog = ""
    cmd = "/usr/sbin/cluster/samples/swap_adapter_complete"
    notify = ""
    pre = ""
    post = ""
    recv = ""
    count = 0

HACMPevent:
    name = "network_up"
    desc = "Network up event happens."
    setno = 0
    msgno = 0
    catalog = ""
    cmd = "/usr/sbin/cluster/samples/network_up"
    notify = ""
    pre = ""
    post = ""
    recv = ""
    count = 0
```

The event you choose to modify with the Event Customization Tool is copied from its original location in /usr/sbin/cluster/events into the /usr/HACMP_ANSS/script directory. The copied event script has its name prefixed by CMD_

The tool will also ask you whether you want to configure a pre, post or recovery event for this event. You can choose one, some, all or none. Depending on your choice(s), the tool will copy one or more shell templates into the

/usr/HACMP_ANSS/script directory. These templates will have the same name as the event but will be prefixed by PRE_, POS_, or REC_, appropriate to your choice.

10.5 Event Customization Tool Example

To start the tool, issue the following command:

```
# /usr/HACMP_ANSS/tools/EVENT_TOOL/event_select
```

After replying to the questions asked, you will see the following panel:

```
++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++
+                                                                    +
+               MENU: Modifying the events                          +
+                                                                    +
+                Choose one option at a time                         +
+          You can choose different events successively              +
+                                                                    +
+                Enter: end (when you have finished)                 +
+                                                                    +
++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++

 1) end                                17) node_down_local
 2) swap_adapter                       18) node_down_local_complete
 3) swap_adapter_complete              19) node_down_remote
 4) network_up                         20) node_down_remote_complete
 5) network_down                       21) node_up_local
 6) network_up_complete                22) node_up_local_complete
 7) network_down_complete              23) node_up_remote
 8) node_up                            24) node_up_remote_complete
 9) node_down                          25) release_service_addr
10) node_up_complete                   26) release_takeover_addr
11) node_down_complete                 27) release_vg_fs
12) join_standby                       28) start_server
13) fail_standby                       29) stop_server
14) acquire_service_addr               30) unstable_too_long
15) acquire_takeover_addr              31) config_too_long
16) get_disk_vg_fs                     32) event_error
Which event would you like to modify: 19
```

The tool will create the necessary templates and also create the corresponding event notification script. Suppose, for example, you chose the following two events:

- node_down_remote
- node_up_remote

For each event you have chosen, the tool will ask you whether you would like to add a PRE, POS or REC event with the aid of the following menu:

```
You have selected: 19 node_down_remote
Do you want to configure the PRE, POS and REC events ?
Choose one option at a time, run as many times as desired
Enter end or 4 to exit


You cannot use this procedure to delete events from the ODM
To do this you will have to use smit
1) PRE event
2) POST event
3) RECOVERY event
4) end
enter your choice ?
```

We will choose PRE and POST events for node_down_remote and a PRE event for node_up_remote.

10.5.1 Looking at the ODM

You can see below how the HACMPevent objects have been modified:

```
HACMPevent:
    name = "swap_adapter"
    desc = "Swap adapter event happens. Swapping adapter."
    setno = 0
    msgno = 0
    catalog = ""
    cmd = "/usr/sbin/cluster/events/swap_adapter"
    notify = ""
    pre = ""
    post = ""
    recv = ""
    count = 0
    .
    .
    .
HACMPevent:
    name = "node_down_remote"
    desc = "Script run when it is a remote node which is leaving the cluster."
    setno = 0
    msgno = 0
    catalog = ""
    cmd = "/usr/HACMP_ANSS/script/CMD_node_down_remote"
    notify = "/usr/HACMP_ANSS/script/event_NOTIFICATION"
    pre = "/usr/HACMP_ANSS/script/PRE_node_down_remote"
    post = "/usr/HACMP_ANSS/script/POS_node_down_remote"
    recv = ""
    count = 0
    .
    .
    .
HACMPevent:
    name = "node_up_remote"
    desc = "Script run when it is a remote node which is joining the cluster."
    setno = 0
    msgno = 0
    catalog = ""
    cmd = "/usr/HACMP_ANSS/script/CMD_node_up_remote"
    notify = "/usr/HACMP_ANSS/script/event_NOTIFICATION"
    pre = "/usr/HACMP_ANSS/script/PRE_node_up_remote"
    post = ""
    recv = ""
    count = 0
```

A list of the shell scripts the tool will have created in the script subdirectory is given below. The scripts are copies of the standard HACMP scripts, put into this alternate location, so future PTF updates to the HACMP scripts will not immediately overwrite any customizations. If you wish, you can modify or customize them so that the event behaves as you require for your specific cluster configuration.

```
CMD_node_up_remote
CMD_node_down_remote
```

The templates for the PRE (before), POS (after) and REC (recovery) are also created, where they are requested. For the above example, a PRE event was requested for the node_up_remote event, and PRE and POS events were requested for the node_down_remote event, so the following files are created:

```
PRE_node_up_remote
PRE_node_down_remote
POS_node_down_remote
```

Also, you can see that the event_NOTIFICATION script is automatically identified as an event notification customization, for any event chosen with the tool.

You can also look at the ODM entries for the HACMP events by entering smit hacmp, and selecting the following options:

```
Manage Node Environment
    Change/Show Cluster Events
```

Selecting, for example, our local node and the node_down_remote event results in the following panel:

```
                        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/HACMP_ANSS/script>

  Notify Command                                  [/usr/HACMP_ANSS/script>
  Pre-event Command                               [/usr/HACMP_ANSS/script>
  Post-event Command                              [/usr/HACMP_ANSS/script>
  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
```

If you pressed the right arrow key in the appropriate fields, you could see the locations of the event customization scripts.

10.5.2 Customizing the Scripts

We will customize the PRE_node_up_remote script to send mail about the event to our main system administrator, and also to send out an immediate message to all users. The message warns those users from the node goofy that it is coming back online, and that they should logoff and wait a few minutes before logging back in.

The customized script is shown below:

```
#!/bin/ksh
# Program   : PRE_node_up_remote
# Role      : run before the event
# Arguments :   $1 = event name
#               and the parameters passed in
# Written   : Wed Dec 13 16:50:41 CST 1995
# Modified  :
 . /usr/HACMP_ANSS/tools/tool_var
STATUS=0
(print "\n=PRE-EVENT==============$(date)"
print "on : $(hostname) "
print "BEFORE : $1"
shift
print "Input Parameters: $*" )  >> $LOG
#####################################################################
# Enter your customizing code here

mail -s "Event Alert" sysadm@theboss.company.com << END
Node goofy is about to re-enter the cluster. Users will be
migrated back from node mickey.
END

wall "Machine goofy has been recovered and is coming on-line.
There will be a short interruption for users of machine goofy.
Please logoff your application now.
You will be able to login to your application again within 5 minutes."
sleep 10

#################### END OF CUSTOMIZATION ########################
return $STATUS
```

In a similar way, you can customize the other PRE and POST event scripts.

10.6 Synchronizing the Node Environment

When you have finished doing your customizations, be sure to synchronize the node environment from the node where you have been working to all the others, before you restart the cluster.

To do this, enter the SMIT fastpath command `smit hacmp` and select the following options:

```
Manage Node Environment
    Sync Node Environment
```

10.6.1 Logging the Events

To check that your customized event scripts are functioning correctly, you can output debug comments into the `/var/HACMP_ANSS/log/hacmp.eventlog` file. This file should be checked periodically. The messages sent into it are put there by the event_NOTIFICATION script, which also allows the possibility of sending mail messages if required.

An example of the output sent by event_NOTIFICATION into `/var/HACMP_ANSS/log/hacmp.eventlog` is shown below:

```
=ODM_EVENT====================Wed Dec 13 16:43:27 CST 1995
Modification of object ++ node_up_remote ++ in HACMPevent
adding customized procedures  PRE
return code = 0

=ODM_EVENT====================Wed Dec 13 16:50:43 CST 1995
Modification of object ++ node_down_remote ++ in HACMPevent
adding customized procedures  PRE POS
return code = 0

=NOTIFICATION===============Mon Dec 18 14:21:11 CST 1995
on: mickey

=PRE-EVENT===============Mon Dec 18 14:21:12 CST 1995
on : mickey
BEFORE : node_down_remote
Input Parameters: goofy graceful
START: node_down_remote
arguments: goofy graceful

=POST-EVENT===============Mon Dec 18 14:21:12 CST 1995
on : mickey
AFTER : node_down_remote
return code : 0

=NOTIFICATION===============Mon Dec 18 14:21:13 CST 1995
on: mickey
OUTPUT: node_down_remote
return code : 0
```

10.7 Testing the Event Customizations

Make sure that you have access to all of the cluster nodes, and that there are no clients connected or using the application(s). Here are some suggested tests:

1. Start HACMP on the nodes and try to provoke a few failures. If you have no subtle solutions, powering off is generally a good way of provoking a failover. Disconnecting the network adapter cable will generate network events. Powering off external disks will create LVM errors.
2. You should NEVER disconnect the SCSI cables because you would risk seriously damaging the disks.
3. Test your application restart on the backup machine.

Chapter 11. Cluster Documentation

This step is carried out after you have configured all of the cluster nodes and your tests have been carried out. The output is a snapshot of your cluster containing:

- Cluster configuration
- Details of any HACMP customization you have carried out
- Scripts you have written
- System files used/modified by HACMP

You have three options for printing the output:

1. ASCII file which can be printed out under AIX
2. Bookmaster file for printing out on a VM host
3. PostScript file produced by the troff command

The report for each machine is called /tmp/HACMPdossier-<hostname>-vm or /tmp/HACMPdossier-<hostname>-ascii or /tmp/HACMPdossier-<hostname>-ps depending upon whether you replied **vm** or **ascii** or **postscript** when you ran the documentation tool.

Nothing prevents you from doing all of them. Obviously, you would need to run the tool multiple times.

An example report, from the doc_dossier tool, is provided in Part 1, "Cluster Documentation Tool Report" on page 137.

11.1 Generating your Cluster Documentation

On one of your cluster nodes, issue the following command:

```
# /usr/HACMP_ANSS/tools/DOC_TOOL/doc_dossier
```

Once the command has executed, a menu will appear on the screen. You should select option 4) Save the output on a UNIX diskette. If you don't have a formatted diskette, choose option 3 first.

Take the diskette produced by the first step to the second cluster node, and restore it by issuing the following command:

```
# tar -xvf/dev/fd0
```

Once you have run doc_dossier on this machine, and returned to the menu, choose option 4.

The diskette now contains the configurations of the two machines.

11.2 Printing the Report on a UNIX System

If you have access to a printer from your system, then you can print the ASCII or PostScript file directly as an option at the completion of a running of the doc_dossier script, or by using the qprt or lp command on the resulting report files left in the /tmp directory.

1. Restore the diskette you have just created using the tar command, if the files are not already on your machine.
2. Print the files named HACMPdossier-<hostname>-ascii or HACMPdossier-<hostname>-ps as appropriate.

11.3 Printing the Report on a VM System

To print the report on a VM system, you will first need a RISC System/6000 connected to that system.

1. Restore the UNIX diskette you created earlier, if necessary.
2. Transfer the files named HACMPdossier-<hostname>-vm to the VM host. You can transfer them using your favorite file transfer program, such as e789 or ftp. Give the VM files a filetype of SCRIPT on the VM host system. If you are using e789 to transfer the files, you will need to set the attributes variable format and record length = 132.
3. To create the LIST3820 file, use the appropriate VM printing command for your system, using at least the twopass option.

   You could also use the dcf command script.

Appendix A. Qualified Hardware for HACMP

The following is the most current copy, as of the writing of this book, of a document called *HAMATRIX*. This document lists the disk adapters, disks, cables, network adapters, and CPU models that are qualified for use with HACMP. By qualified, this means that the device has been tested by IBM, with HACMP, so the user can have a high degree of confidence that there will not be mysterious errors with the device that cannot be fixed.

The HAMATRIX document is maintained on an IBM tools disk called *MKTTOOLS*. If you are planning on implementing HACMP, or are considering adding new hardware to an existing cluster, contact your IBM representative to receive the latest version of this document.

A.1 The HAMATRIX Document

DISK STORAGE MEDIA, PROCESSORS AND ADAPTERS
QUALIFIED FOR USE WITH HACMP FOR AIX

Document Version 4.1A
8/17/95

This document designates which hardware has been qualified for use with HACMP for AIX (herafter referred to as HACMP). The designated hardware should only be used on an appropriate RISC System/6000 Platform or 9076 Scalable POWERParallel Platform (SP/2). Please refer to the processor documentation to be sure that appropriate hardware is obtained.

This document contains the following information:

- The main body of the document and Appendix A contain the disk adapters, disk enclosures and associated cabling;
- Appendix B contains other hardware, e.g. processors and network adapters.

The document is intended to convey information pertinent to HACMP support so cabling methods and hardware features unrelated to HACMP are not shown. If a piece of hardware is not listed it should be assumed that the hardware is not supported by HACMP.

The following are the major changes since the last version of this matrix:

- Serial Storage Architecture (SSA) supported on HACMP Version 3.1.1
- Enhanced SCSI-2 Fast/Wide Adapter/A (FC 2412) supported on HACMP Version 3.1.1
- Target Mode on SCSI-2 Fast/Wide Adapters (FC 2412 and FC 2416) supported on HACMP Version 3.1.1
- IBM RISC System/6000 7013 Model 591, 7015 Model R21 and 7015 Model R3U

The disk storage portions of the document contain brief descriptions of many of the disk drive adapters, disk enclosures and associated cabling in tabular form. These tables are grouped as follows and unless specifically noted otherwise, the hardware in one group can not be used with hardware in another group:

- SCSI-2 Differential Device Support
- Serial Device Support

One of the columns in the disk tables is titled ″HACMP RIse″ and contains two subheadings:

- Non-concurrent disk access, denoted by an NC in the column heading (Modes 1 and 2)
- Concurrent disk access, denoted by a CC in the column heading (Mode 3)

Under each subheading in the disk tables is noted the release of HACMP in which the hardware was first supported for that configuration. The following conventions were used for this data:

- If the specified release is prior to the current release, then the hardware is still supported unless noted otherwise.
- If the column has a TBD in it then no commitment has been made to support the hardware; the hardware might or might not be supported in the future.
- If the column has an N/A in it then there are no plans to support the hardware.

Attachment A contains the SCSI-1 SE and SCSI-2 SE device support. Existing HACMP configurations using SCSI SE devices continue to be supported. New HACMP installations must use SCSI-2 differential or serial devices due to the unavailability of the PTT cables.

If you have further questions about disk cabling you can also consult the following information:

- RISC System/6000, System Overview and Planning, Chapter 7: Cables and Cabling (GC23-2406)
- A copy of the SCSI cabling portion of publication GC23-2406 can be found on MKTTOOLS(RS6CABLE)
- A pictorial view of some of SCSI cabling for HACMP is available in MKTTOOLS(HASCSI6)

(The proper hardware documents take precedence over the hardware information contained in these tables and should be used to resolve any conflicts.)

SCSI-2 DIFFERENTIAL DEVICE SUPPORT

The following conventions are used in this section:

- All 16 bit adapters and enclosures have an * next to their feature codes. All 16 bit cables or 8 bit to 16 bit cables have an * next to their feature codes. The 16 bit implementation is generally known as SCSI Fast/Wide.
- Enclosures which can be cabled with either 16 bit or a combination of 8 bit and 16 bit cables have @ next to their feature codes.
- All 8 bit adapters, enclosures and cables have no indication next to their feature codes.

ADAPTERS

Feature (FRU #)	MBPS	Maximum Cable Length	HACMP Rlse NC	HACMP Rlse CC	Notes
2412*	20	25 m	3.1.1	3.1.1	(2,3,5,6,7,8,9)
2416* (65G7315)	20	25 m	2.1	2.1	(2,3,5,6,7,8,9)
2420 (43G0176)	10	19 m	1.2	1.2	(1,2,3,4)

Notes:

1 - Eight external SCSI IDs and eight LUNs are available on these buses. In an HACMP environment two or more of the addresses are used for hosts so the bus can have up to a maximum of six other devices (subject to cabling length and device constraints).
2 - Only SCSI-2 differential devices can be attached to a SCSI-2 differential adapter.
3 - Cable length is measured from end to end and includes the cabling which is within any attached subsystems. Exception: For the 7135, no internal SCSI-2 SE cabling is included.
4 - In HACMP configurations the differential terminating resistors U8 and U26 must be removed from the 2420 adapter; these resistors are located next to the external SCSI bus connector on the adapter card.
5 - 2412 and 2416 adapter can execute in either 8 bit or 16 bit mode; a SMIT option exists to set the adapter to the desired width. All the devices on the bus must of the same type.
6 - HACMP does not support target mode SCSI on the 2412 or the 2416 adapter prior to HACMP Version 3.1.1; on HACMP Version 3.1.1 APAR IX52772 is required.
7 - In HACMP configurations the three built-in differential terminating resistors (labelled RN1, RN2 and RN3) must be removed from the 2412 and 2416 adapters.
8 - In HACMP Version 4.1 sixteen external SCSI IDs and 32 LUNs are available on these buses. In an HACMP environment two or more of the addresses are used for hosts so the bus can have up to a maximum of fourteen other devices (subject to cabling length and device constraints). Prior to HACMP Version 4.1 eight external SCSI IDs and eight LUNs are available on these buses. In an HACMP environment two or more of the addresses are used for hosts so the bus can have up to a maximum of six other devices (subject to cabling length and device constraints).
9 - The 2412 and 2416 can not be assigned SCSI IDs 0, 1 or 8 through 15.

ENCLOSURES

Model	# Per Bus	# Dsk Drv	Size GB	Disk Feat	Media Rate MBPS	HACMP Rlse NC	HACMP Rlse CC	Notes
7204-215	4	1	2.0	-	5.22	2.1	N/A	(1)
7204-315*	6	1	2.0	-	5.22	2.1	N/A	(1,8)
7204-317*	14	1	2.2	-	9-12	3.1	N/A	(1,8)
7204-325*	14	1	4.5	-	9-12	3.1	N/A	(1,8)
9334-011	2	4	1.0	2565	3.0	1.2	N/A	(1,4)
	2	4	2.0	2585	5.22	1.2	N/A	(1,4)
9334-501	2	4	1.0	2565	3.0	1.2	N/A	(1,4)
	2	4	2.0	2585	5.22	1.2	N/A	(1,4)
7134-010*	1	16	2.0	2821	5.22	2.1	N/A	(1,5)
	1	16	2.2	2712	9-12	3.1	N/A	(1,5)
	1	16	4.5	2714	9-12	3.1	N/A	(1,5)
7135-010	-	12	2.0	2720	5.22	N/A	N/A	(1)
7135-110@	2	30	1.3	2715	5.22	(7)	(7)	(1,2,3,7)
	2	30	2.0	2725	5.22	(7)	(7)	(1,2,3,7)
	2	30	2.2	2825	9-12	(7)	(7)	(1,2,3,7)
	2	30	4.5	2845	9-12	(7)	(7)	(1,2,3,7)
7135-210@	2	30	1.3	2715	5.22	4.1	4.1	(1,2,3)
	2	30	2.0	2725	5.22	4.1	4.1	(1,2,3)
	2	30	2.2	2825	9-12	4.1	4.1	(1,2,3)
	2	30	4.5	2845	9-12	4.1	4.1	(1,2,3)
3514-212@	2	8	1.0	1011	5-6	2.1	2.1	(1,6)
3514-213@	2	8	2.0	1008	5.22	2.1	2.1	(1,6)
7137-412@	2	8	1.0	1020	5.22	2.1	2.1	(1,6)
7137-413@	2	8	2.0	1030	9-12	2.1	2.1	(1,6)
7137-414@	2	8	4.4	1040	9-12	2.1	2.1	(1,6)
7137-512@	2	8	1.0	1020	5.22	2.1	2.1	(1,6)
7137-513@	2	8	2.0	1030	9-12	2.1	2.1	(1,6)
7137-514@	2	8	4.4	1040	9-12	2.1	2.1	(1,6)

Notes:

1 - All SCSI-2 Differential devices use one bus address per disk except the 7135, 3514 and the 7137 which use one address per controller. All devices on the same bus must be of the same type unless stated otherwise.

2 - For maximum availability the 7135 array should be configured with two controllers. HACMP supports RAIDs 1, 3 and 5. The external interface for the 7135 is SCSI-2 differential; however, internally the disk drives are SCSI-2 SE.

3 - The specified disk feature provides a full bank of five disks. Disks in the 7135 array are normally configured in banks of 5 disks each, for a total capacity of 30 disks.

4 - 9334-011 and 9334-501 enclosures can be daisy chained with up to two enclosures and six disk drives on a SCSI bus. No tape drives are permitted.

5 - With two hosts the 7134-010 without an internal expansion unit can support up to eight drives on one bus. With an internal expansion unit the maximum number of drives with two hosts and one bus is fourteen. With an internal expansion unit the maximum number of drives with two hosts and two buses is sixteen.

6 - Even though the 3514 and 7137 are RAID devices, they have single

points of failure in the SCSI bus and in the controller. If this is unacceptable, one or more additional enclosures with LVM mirroring are required; a total of three enclosures with quorum provides the "highest" availability. Concurrent access mode (HACMP Mode 3) will not support mirroring on SCSI devices so the single points of failure noted above would exist in this configuration.

7 - HACMP Version 4.1 does not support the 7135-110. The 7135-110 is supported in HACMP Version 2.1 and later releases, up to but not including HACMP Version 4.1.

8 - 7204 Models 315, 317 and 325 can be used on the same SCSI-2 differential bus.

CABLES

Feature (Part #)	Attachd From	Attachd To	Len (m)	Notes
CONFIGURED ON SERVERS WITH 8 BIT WIDE ADAPTER				
2422 (52G7348)	Adapter (2420)	9334 cable, 3514 cable*, 7137 cable*, 7204-215 cable, terminator, 2423	.765	Y-cable: o base to adapter; o 8 bit long leg to - 9334 cable, - 3514 cable, - 7137 cable or - 7204-215 cable; o 8 bit short leg is - terminated or - connected to a 2423 cable to add additional processors (>2 processors) to a shared differential 8-bit bus
N/A (52G7350)	Y-cable (2422)	self	0	Terminator, 8 bit, included when the Y-cable is ordered.
2423 (52G7349)	Y-cable (2422, 2427)	Y-cable (2422, 2427) on other system	2.5	Cable can be used to attach a third and fourth system to a shared differential 8 bit bus.
CONFIGURED ON SERVERS WITH 16 BIT WIDE ADAPTER				
2427* (52G4349)	Adapter (2412*, 2416*)	9334 cable, 7204-215 cable, 2424*/2425*, terminator*	.765	Y-cable: o 16 bit base to adapter; o 8-bit long leg to - 9334 cable or - 7204-215 cable; o 8-bit short leg is - terminated or - connected to a 2423 cable to add additional processors (>2 processors)
2426* (52G4234)	Adapter (2412*, 2416*)	7204-3XX cable*, 3514 cable*, 7137 cable*, 7134-010 cable*, 2424*, 2425*, terminator*	.94	Y-cable: o 16 bit base to adapter; o 16-bit long leg to - 7204-3XX cable, - 3514 cable, - 7137 cable or - 7134-010 cable; o 16-bit short leg is terminated or is connected to a 2424 or 2425 cable to add additional processors (>2 processors)
2426* (52G4234)	Adapter (2412*)	7135-210 cable*, 2424*, 2425*, terminator*	.94	Y-cable: o base to adapter; o 16-bit long leg to - 7135-210 cable; o 16-bit short leg is terminated

				or is connected to a 2424 or 2425 cable to add additional processors (>2 processors)
2426* (52G4234)	Adapter (2416*)	7135-110 cable*, 2424*, 2425*, terminator*	.94	Y-cable: o base to adapter; o 16-bit long leg to - 7135-110 cable; o 16-bit short leg is terminated or is connected to a 2424 or 2425 cable to add additional processors (>2 processors)
N/A* (61G8324)	Y-cable (2426*)	self	0	Terminator, 16-bit, included when the Y-cable is ordered.
N/A (52G7350)	Y-cable (2427*)	self	0	Terminator, 8 bit, included when the Y-cable is ordered.
2424*/2425*	Y-cable (2426*)	Y-cable (2426*) on other system	 .6 2.5	Cable can be used to attach a third and fourth system to a shared differential 16 bit bus. 2424 (52G4291) 2425 (52G4233)
CONFIGURED ON 7204-215 **********				
2854/2921	Y-cable (2422, 2427*)	7204-215	 0.6 4.75	Needed on 7204-215 at each end of the shared unit. 2854 (87G1358) 2921 (67G0593)
2848 (74G8511)	7204-215	7204-215	2.0	Used between 7204-215's on the shared string.
CONFIGURED ON 7204-315, 7204-317, 7204-325 **********				
2845*/2846*	Y-cable (2426*)	7204-315*, 7204-317*, 7204-325*	0.6 2.5	2845 (52G4291) 2846 (52G4233) Needed on 7204-3XX at each end of the shared unit.
2845*/2846*	7204-315* 7204-317* 7204-325*	7204-315*, 7204-317*, 7204-325*	0.6 2.5	2845 (52G4291) 2846 (52G4233) Used between 7204-3XX's on the shared string.
CONFIGURED ON 9334-011 **********				
2921/2923	Y-cable (2422, 2427*)	9334-011	 4.75 8.0	Needed on 9334-011 at each end of the shared unit. 2921 (67G0593) 2923 (95X2494) To conform to the cable length limit, the 8.0 meter cable must be paired with the 4.75 meter cable.
2925 (95X2492)	9334-011	9334-011	2.0	Allows daisy chaining of two 9334-011 enclosures

CONFIGURED ON 9334-501
****

2931/2937	Y-cable (2422, 2427*)	9334-501		Needed on 9334-501 at each end of the shared unit.
			1.48	2931 (7OF9188)
			2.38	2933 (45G2858)
			4.75	2935 (67GO566)
			8.0	2937 (67GO562)
				To conform to the cable length limit, the 8.0 meter cable must be paired with a shorter cable.
2939 (95X2498)	9334-501	9334-501	2.0	Allows daisy chaining of two 9334-501 enclosures

CONFIGURED ON 7134-010
****

2902-2918*	Y-cable (2426*)	7134-010*		Needed on 7134-010 at each end of the shared unit.
			2.4	2902 (88G5750)
			4.5	2905 (88G5749)
			12.0	2912 (88G5747)
			14.0	2914 (88G5748)
			18.0	2918 (88G5746)

CONFIGURED ON 7135-110 AND 7135-210
****

2919 (61G8323)	Y-cable (2422)	7135 cable*	0	Cable interposer; connects 8 bit Y-cable to 16 bit 29XX cable for 7135
2901*-14*	2919, Y-cable (2426*)	7135@		Connects 7135 array controller to an interposer (2919) or to a 16 bit Y-cable
			0.6	2901 (67G1259)
			2.4	2902 (67G1260)
			4.5	2905 (67G1261)
			12	2912 (67G1262)
			14	2914 (67G1263)
			18	2918 (67G1264)
				To conform to the cable length limit, the 12m, 14m and 18m cables must be paired with shorter cables.

CONFIGURED ON 3514
****

2002*	Y-cable (2422*)	3514@	4.0	Needed on 3514 at each end of the shared unit (8-bit to 16-bit cable)
2014*	Y-cable (2426*)	3514@	4.0	Needed on 3514 at each end of the shared unit
3001*	3514*	3514@	2.0	Allows daisy chaining of two 3514 units

CONFIGURED ON 7137
****

2002*	Y-cable (2422*)	7137@	4.0	Needed on 7137 at each end of the shared unit (8-bit to 16-bit cable)
2014*	Y-cable	7137@	4.0	Needed on 7137 at each end of

	(2426*)			the shared unit
3001*	7137*	7137@	2.0	Allows daisy chaining of two 7137 units

Notes:
---

1 - After configuring a SCSI-2 differential bus for the HACMP environment , use the following checklist to validate the configuration:
  - At least two and no more than four processors are attached to the bus.
  - Only SCSI-2 differential cables, adapters and devices were used.
  - A Y-cable is attached to each processor on the bus.
  - The bus must have a terminator on the short leg of each Y-cable which is at the end of the bus (total of 2 terminators per bus).
  - 8 bit wide and 16 bit wide enclosures can not be used on the same bus.
  - You must not exceed maximum SCSI-2 differential bus lengths, including the cabling within enclosure cabinets. Cable lengths within enclosure cabinets are:
    - 7204-215 nil
    - 7204-315 nil
    - 7204-317 nil
    - 7204-325 nil
    - 9334-011 3.1 meters
    - 9334-501 2.66 meters
    - 7134-010 3.0 meters/bus
    - 7135-110 0.66 meters/controller
    - 7135-210 0.66 meters/controller
    - 3514-2XX 1.0 meters
    - 7137-XXX 0.2 meters

  The publication "Common Diagnostics and Service Guide" (SA23-2687) contains additional information about cabling.

2 - For a given cable, any item listed in the "Attachd From" column can be connected to any item in the "Attachd To" column. Y-cables do not follow this rule; they have three legs and the above tables show what connects to each of the legs.

3 - The configurations in this table assume that processors are at the two ends of the bus (just prior to each terminator) and all the storage devices are connected to the bus between the processors.

4 - The recommended 7135 configuration for HACMP is:
  - Two controllers on the 7135, each controller on a separate SCSI-2 differential bus
  - Each controller is attached to every processor in the cluster.

  This yields two different SCSI-2 differential buses, each bus is connected to one controller and to every processor in the cluster. The Disk Array Manager software in the processors manages access to the different controllers and will switch controllers if one of the controller fails; this occurs independently of HACMP.

5 - SCSI buses can not include non-disk devices (i.e. tape, CD ROM).

SERIAL DEVICE SUPPORT

= =

ADAPTERS

Feature (FRU #)	MBPS	HACMP Rlse NC	CC	Notes
6210 (52G1071)	8	1.1	1.2	(1,2,3)
6211 (00G3357)	8	1.1	1.2	(1,2,3)
6212 (67G1755)	8	1.2	1.2	(1,2,3)

Notes:

1 - Only serial devices can be attached to a serial adapter.
2 - For serial adapters the maximum cable length is measured from the adapter to the subsystem controller. The cabling which might be within a subsystem is not included.
3 - Serial adapters contain four serial link connectors to allow the attachment of up to four serial subsystems (e.g. four 9333's). Data transfer rates on the microchannel side of the adapter are:
    6210 - 40 MBPS, used for 9333 Model 010 or Model 500
    6211 - 80 MBPS, used for 9333 Model 010 or Model 500
    6212 - 40 or 80 MBPS, used for 9333 Model 011, Model 501, Model 010 or Model 500

ENCLOSURES

Model	# Dsk Drv	Size -GB-	Disk Feat	Media Rate MBPS	HACMP Rlse NC	CC	Notes
9333-010	4	.857	3100	3.0	1.1	1.2	
	4	1.07	3110	3.0	1.1	1.2	
9333-011	4	.857	3100	3.0	1.2	1.2	
	4	1.07	3110	3.0	1.2	1.2	
	4	2.0	3120	5.22	1.2	1.2	
9333-500	4	.857	3100	3.0	1.1	1.2	
	4	1.07	3110	3.0	1.1	1.2	
9333-501	4	.857	3100	3.0	1.2	1.2	
	4	1.07	3110	3.0	1.2	1.2	
	4	2.0	3120	5.22	1.2	1.2	

Notes:

1 - The following table shows the HACMP support for the 9333:

AIX Release	3.2.3E		3.2.4			3.2.5			
HACMP Release	1.2		1.2		2.1	1.2		2.1	
Configuration	NC	CC	NC	CC	NC	CC	NC	CC	NC	CC
9333 010/500	2	2	2	N	2	N	2	N	2	N
PTF #	-	-	-	-	-	-	-	-	-	-

9333	011/501	N	N	2	2	2	2	2	2	4	4
	PTF #	-	-	-	a	-	-	-	b	-	c

N = Not supported
2 = 2-way is supported, if PTF# is not specified then the support is in the base system. Under AIX 3.2.4 Feature codes 4001 and 4002 of the 9333-011 and -501 subsystem are not permitted.
4 = 2-, 3- and 4-way are supported, if PTF# is not specified then the support is in the base system. If either 3- or 4-way is desired then Feature 4001 must be installed on the 9333-011 or -501.
a = U421401 or supersede
b = U425614 or supersede
c = U426577 or supersede

2 - 9333 Models 010 and 500 come standard with two ports connected to one controller card; the controller card controls up to 4 disks inside the enclosure. The ports can be connected to two different hosts using one serial link connector on each host adapter. An upgrade is available to go from a 9333 Model 010 to a 9333 Model 011, or from a 9333 Model 500 to a 9333 Model 501.

3 - 9333 Models 011 and 501 come standard with two ports connected to one controller card; the controller card controls up to 4 disks inside the enclosure. The ports can be connected to two different hosts using one serial link connector on each host adapter. With the 9333 Models 011 or 501, the number of attachable hosts can be expanded by ordering the appropriate expansion features, either to 4 systems (feature 4001) or to 8 systems (features 4001 and 4002).

4 - The data transfer rate for a serial bus is 8 MB/sec.

```
CABLES
------

Notes:
------
1 - There are no special cabling requirements for HACMP for AIX. The
    publication "Common Diagnostics and Service Guide" (SA23-2687)
    contains information about cabling serial buses.
2 - Each 9333 enclosure comes standard with one attachment cable.
    Additional cables need to be ordered to attach it to more than one
    system.
```

SERIAL STORAGE ARCHITECTURE (SSA)

= =

ADAPTERS

Feature (FRU #)	MBPS	HACMP Rlse NC	HACMP Rlse CC	Notes
6214	80	(1)	(1)	(1,2)

Notes:

1 - The 6214 adapter is supported on HACMP Version 3.1.1 only; APAR IX52776 is required.
2 - Only two 6214 adapters can be put into a single SSA loop; one in each processor in the cluster.

ENCLOSURES

Model	# Dsk Drv	Size -GB-	Disk Feat	Media Rate MBPS	HACMP Rlse NC	HACMP Rlse CC	Notes
7133-010	16	1.1	31XX	35	(1)	(1)	(1,2,3)
	16	2.2	32XX	35	(1)	(1)	(1,2,3)
	16	4.5	34XX	35	(1)	(1)	(1,2,3)
7133-500	16	1.1	31XX	35	(1)	(1)	(1,2,3)
	16	2.2	32XX	35	(1)	(1)	(1,2,3)
	16	4.5	34XX	35	(1)	(1)	(1,2,3)

Notes:

1 - The 7133-010 and 7133-500 are supported on HACMP Version 3.1.1 only; APAR IX52776 is required.
2 - The disk features are YYXX where YY is as shown in the table above and XX is 01, 08 or 16 for one, eight or sixteen
3 - Up to 96 disks can be supported in a single SSA loop.

CABLES

Notes:

1 - There are no special cabling requirements for HACMP. The publication "Common Diagnostics and Service Guide" (SA23-2687) contains information about cabling.

ATTACHMENT A

Attachment A contains the SCSI-1 SE and SCSI-2 SE device support. Existing HACMP configurations using SCSI SE devices continue to be supported. New HACMP installations must use SCSI-2 differential or serial devices due to the unavailability of the PTT cables.

The SCSI SE PTT cables (FC 2914 and FC 2915) are available via an RPQ but only with prior Austin lab approval of the specific configurations. Two of these cables are required for a minimum HACMP configuration.

None of the equipment in this attachment can be configured in a new HACMP installation.

SCSI-1 SE AND SCSI-2 SE DEVICE SUPPORT

= =

ADAPTERS

Feature (FRU #)	MBPS	TYP	Maximum Cable Length	HACMP Rlse NC	HACMP Rlse CC	Notes
2835 (31G9729)	4	1	6 m	1.1	N/A	(1,2,3,4)
2410 (52G5484 52G7509)	10	2	4.75 m	1.2	N/A	(1,2,3,5)
2415	20	2	note 7	N/A	N/A	(1,2,3,6,7)

Notes:
------

1 - Eight external device addresses are available on these buses. In an HACMP environment two of the addresses are used for hosts so the bus can have up to six other devices (subject to cabling length constraints).

2 - Only SCSI SE devices can be attached to a SCSI SE adapter.

3 - Cable length is measured from one end of the bus to the other and includes the cabling which is within any attached disk subsystem enclosures.

4 - In an HACMP environment the 2835 adapter can only be used with SCSI-1 SE disk enclosures. Minimum assembly numbers which can be used for an HACMP configuration is part #31G9722 and Field Replaceable Unit (FRU) #31G9729. For HACMP configurations the 50 position card edge terminator must be removed, and the jumper J1 must be removed. The removed jumper can be moved over and attached to only one row of pins for storage, the row furthest from the the external SCSI connector.

5 - In an HACMP environment the 2410 adapter can only be used with the 7203 and/or 7204 enclosures utilizing the 1 GB SCSI-2 SE disk, (7203-001 with feature 2320 or 7204-001). For HACMP configurations the 50 position card edge terminator must be removed, and the jumper P3 must be removed. The removed jumper can be moved over and attached to only one row of pins for storage, the row furthest from

the external SCSI connector.
6 - This adapters can execute in either 8 bit or 16 bit mode; a SMIT option exists to set the adapter to the desired width. All the devices on the bus must of the same type.
7 - Maximum cable length varies with the configuration:
- 6m when attached to 9334-500
- 3m what attached to anything else.

ENCLOSURES

Model	T Y P	# Per Bus	# Dsk Drv	Size -GB-	Disk Feat	Trans. Rate MBPS Media	Bus	HACMP Rlse NC	CC	Notes
7203-001	1	4	1	.355	2300	1.87	4	1.1	N/A	(5)
	1	4	1	.670	2310	1.87	4	1.1	N/A	(5)
	2	2	1	1.0	2320	5.0	5	1.2	N/A	(3,5)
7204-320	1	5	1	.320	-	2.0	4	1.1	N/A	
7204-001	2	2	1	1.0	-	3.0	5	1.2	N/A	(3,5)
7204-010	2	-	1	1.0	-	3.0	5	N/A	N/A	
9334-010	1	-	4	.670	2510	1.87	4	N/A	N/A	(1)
	1	-	4	.857	2530	3.0	4	N/A	N/A	(1)
	1	-	4	1.37	2570	4.5	5	N/A	N/A	(1)
	2	-	4	2.0	2580	5.22	10	N/A	N/A	(1)
	2	-	3+1	2.4	2590	3.0	10	N/A	N/A	(1,4)
	2			1.0	2555	3.0	10	N/A	N/A	(1,4)
9334-500	1	1	4	.670	2510	1.87	4	1.1	N/A	(6)
	1	1	4	.857	2530	3.0	4	1.1	N/A	(6)
	1	1	4	1.37	2570	4.5	5	1.2	N/A	(2,6)
	2	-	4	2.0	2580	5.22	10	N/A	N/A	
	2	-	3+1	2.4	2590	3.0	10	N/A	N/A	(4)
	2			1.0	2555	3.0	10	N/A	N/A	(4)

Notes:

1 - The internal cabling of the 9334-010 makes it unsuitable for sharing between systems. Therefore it is not supported by HACMP. Only the 9334-500 is supported, with the features as noted in the table above.
2 - Disk fencing must not be enabled in an HACMP environment unless the fix documented in the HACMP Version 1.2 Release Notes is applied.
3 - For use with HACMP in a twin-tailed environment, 1 GB disks for the 7203 and 7204 enclosures (7203-001 with feature 2320, 7204-001) are only tested and supported using the SCSI-2 SE adapter (feature 2410).
4 - The 2590 which uses two bus addresses is two 1.2 GB disks within a single package. The 2555 drive is available only as the fourth drive within a 9334 which contains 3 2590's.
5 - The limitation in the table under "# Per Bus" is not a cabling limitation but a testing limitation and only the specified number of devices is supported on the bus. (Cable limitations allow one more device to be connected than is shown.)
6 - 9334-500 in an HACMP environment is supported only on the 2835 adapter.

CABLES

Feature (Part #)	Type	Attachd From	Attachd To	Len (m)	Notes
3130 (31F4222)	SCSI-1/2 SE	7203, 7204	7203, 7204	0.66	Device-to-Device cable. Used between devices in a shared string.
2915 (00G0959)	SCSI-1 SE	Adapter (2835)	7203, 7204	1.57	Passthru terminator (PTT) cable, withdrawn from marketing. See note #4.
2915 (70F9171)	SCSI-1 SE	Adapter (2835)	9334-500	1.48	Passthru terminator (PTT) cable, withdrawn from marketing. See note #4.
2914 (51G8568)	SCSI-2 SE	Adapter (2410)	7203, 7204	1.57	Passthru terminator (PTT) cable, withdrawn from marketing. See note #4.

Notes:

1 - After configuring a SCSI SE bus for the HACMP environment, use the following checklist to validate the configuration:
  - Two processors must be attached to the bus.
  - Only SCSI SE cables, adapters and enclosures can be used.
  - A shared SCSI SE bus requires two PTT cables, one attached to each adapter.
  - You must not exceed maximum SCSI SE bus lengths, including the cabling within enclosure cabinets. The SCSI SE maximum bus cable lengths are:
    - SCSI-1 SE 6 meters
    - SCSI-2 SE 4.75 meters

    Cable lengths within enclosure cabinets:
    - 7203 nil
    - 7204 nil
    - 9334-010 not supported by HACMP
    - 9334-500 2.66 meters

  The publication "Common Diagnostics and Service Guide" (SA23-2687) contains additional information about cabling.
2 - For a given cable, any item listed in the "Attachd From" column can be connected to any item in the "Attachd To" column
3 - SCSI bus can not include non-disk devices (i.e. tape, CD ROM)
4 - The PTT cables are available via an RPQ but only after the Austin lab approves the specific SCSI SE bus configuration(s) involved. FC 2915 is available via RPQ #8A0759; FC 2914 is available via RPQ #8A0758.

ATTACHMENT B
OTHER HARDWARE QUALIFIED WITH HACMP

PROCESSORS

7009-C10	7012-320	7013-52H	7013-58F	7015-97E
7009-C20	7012-34H	7013-520	7013-58H	7015-97F
7011-22W	7012-340	7013-53E	7013-580	7015-970
7011-220	7012-350	7013-53H	7013-59H	7015-98B
7011-23S	7012-355	7013-530	7013-590	7015-98E
7011-23T	7012-36T	7013-540	7013-591	7015-98F
7011-23W	7012-360	7013-55E	7015-R10	7015-980
7011-230	7012-365	7013-55L	7015-R20	7015-99E
7011-25S	7012-37T	7013-55S	7015-R21	7015-99F
7011-25T	7012-370	7013-550	7015-R24	7015-99J
7011-25W	7012-375	7013-56F	7015-930	7015-99K
7011-250	7012-380	7013-560	7015-95E	7015-990
7012-32E	7012-39H	7013-57F	7015-950	
7012-32H	7012-390	7013-570	7015-97B	

Symmetric Multi-Processors
7012-G30, 7013-J30, 7015-R30 and 7015-R3U

9076 Scalable POWERParallel Platforms (SP/2) - supported on HACMP Version 3.1.1 but not HACMP Version 4.1

Asynchronous Communication Adapters

- FC 2930 - 8 Port Async Adapter - EIA-232
- FC 2950 - 8 Port Async Adapter - MIL-STD 188
- FC 2955 - 16 Port Async Adapter - EIA-232
- FC 6400 - 64 Port Async Controller
- FC 8128 - 128 Port Async Controller

Local Area Network (LAN) Communication Adapters

- FC 2402 - Network Terminal Accelerator
  - High performance ethernet adapter permitting up to 256 login sessions when used in conjunction with a 7318 Model S20 Serial Communications Network Server. HACMP supports only the MAC Layer Interface for the adapter, not the HTY functionality.
- FC 2403 - Network Terminal Accelerator
  - High performance ethernet adapter permitting up to 2048 login sessions when used in conjunction with a 7318 Model S20 Serial Communications Network Server HACMP supports only the MAC Layer Interface for the adapter, not the HTY functionality.
- FC 2720 - Fiber Distributed Data Interface Adapter
- FC 2722 - Fiber Distributed Data Interface Dual Ring Upgrade KIT
- FC 1906 - Fiber Channel Adapter/266
- FC 2723 - FDDI / Fiber Dual-Ring Upgrade
- FC 2724 - FDDI - Fiber Single-Ring Adapter
- FC 2725 - FDDI - STP Single-Ring Adapter
- FC 2726 - FDDI - STP Dual-Ring Upgrade
- FC 2970 - Token-Ring High-Performance Network Adapter
- FC 2972 - Auto Token-Ring Lanstreamer 32 MC Adapter

```
   FC 2972 - Auto Token-Ring Lanstreamer 32 MC Adapter
   FC 2980 - Ethernet High-Performance LAN Adapter
   FC 4224 - Ethernet 10BASET Transceiver (Twisted Pair)


RS-232 Serial Network
=====================
   FC 3107 - C10 Serial Port Converter
   FC 3124 - 3.7 Meter Serial to Serial Port Cable
   FC 3125 - 8 Meter Serial to Serial Port Cable


Other Adapters / Subsystems
===========================
   7318-P10  Serial Communications Network Server
            -allows attachment of async devices and parallel printers
             to an Ethernet LAN attached RISC System/6000
             (Most commonly concerned with HACMP configurations when
              used with FC 2402/3 Network Terminal Accelerator)
   7318-S20  Serial Communications Network Server
            -allows attachment of async devices and parallel printers
             to an Ethernet LAN attached RISC System/6000
             (Most commonly concerned with HACMP configurations when
              used with FC 2402/3 Network Terminal Accelerator)
   FC 2860 - Serial Optical Channel Converter
   FC 4018 - High Performance Switch (HPS) Adapter-2
             - supports node fallover on an SP/2



= = = = = = = = = = = = = = = = = = = = = = = = = = = = = = = = = = = = = = = =
                              end of document
= = = = = = = = = = = = = = = = = = = = = = = = = = = = = = = = = = = = = = = =
```

Appendix B. RS232 Serial Connection Cable

In implementing the non-TCP/IP RS232 link between cluster nodes, implementers of HACMP now have at least three choices for the cable:

1. A standard cable for this purpose, marketed by IBM
2. Putting together the correct connection, using a combination of IBM and non-IBM cables and connectors
3. Building a custom cable

B.1 IBM Standard Cable

IBM now markets a special asynchronous communications cable to serve as the HACMP RS232 connection cable. This cable has the correct pinouts configurated to allow the cable to connect a 25-pin RS232 port on one machine to a 25-pin RS232 port on another machine. The newer models of RS/6000 have 25-pin native RS232 ports, where this cable can be used.

If you have an older model, with its 10-pin native RS232 ports, you will have to add a 10-pin to 25-pin converter cable to each end. The part number of this cable is 58F3740.

The standard IBM cable comes in two lengths. The feature numbers are orderable against any RS/6000 CPU model:

- Feature 3124 (Part number 88G4853) - 3.7 meter cable
- Feature 3125 (Part number 88G4854) - 8.0 meter cable

Each of these cables has the null modem pinout connections required to make a direct connection between serial ports.

B.2 Putting together Available Cables and Connectors

If you are going to make up the serial network between the cluster nodes using standard IBM cables you will need the following:

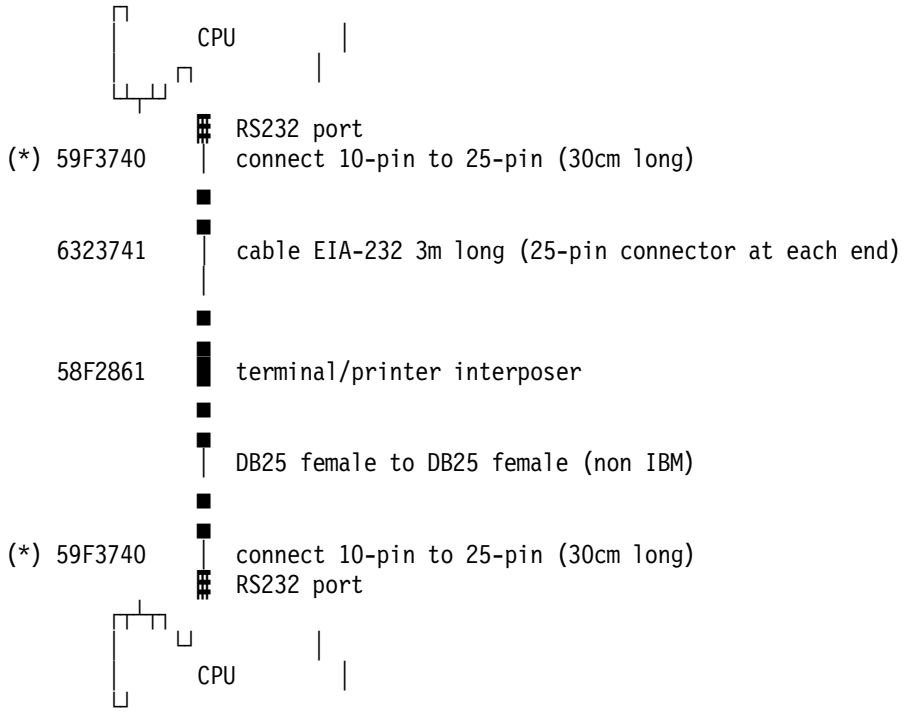

(*) optional if your machine has a 25-pin port

B.3 Making your Own Cable

You can make up your own cable for the serial connection. The wiring scheme is given below:

Female Connector N 1	Signal	Female Connector N 2
1	Shield Ground	shell
2	TxD	3
3	RxD	2
4	RTS	5
5	CTS	4
6,8	DSR,CD	20
7	Signal Ground	7
20	DTR	6,8

Table 1. Wiring scheme for the RS232 connection between nodes

Appendix C. List of AIX Errors

The following is a list of the current AIX errors, capable of being written into the AIX error log. These errors apply to AIX 3.2.5 maintenance level 3251. They are obtained by running the command errpt -t.

Id	Label	Type	CL	Error_Description
00530EA6	DMA_ERR	UNKN	H	UNDETERMINED ERROR
01F2D769	X25_ALERT25	PERM	H	X-25 RESTART REQUEST BY X.25 ADAPTER
0299F00B	FDDI_NOMBUFS	TEMP	S	RESOURCE UNAVAILABLE
03348B46	CXMA_MEM_BD	PERM	S	Can¢t Allocate bd_t Structures
0375DFC2	X25_ALERT9	TEMP	H	X-9 FRAME TYPE W RECEIVED
038F2580	SCSI_ERR7	UNKN	H	UNDETERMINED ERROR
038F3117	MPQP_DSRDRP	TEMP	H	COMMUNICATION PROTOCOL ERROR
03ACD152	NB20	PERM	S	SOFTWARE PROGRAM ERROR
04B1C8C0	VCA_INITZ	TEMP	S	Host independent initialization failed
0502F666	SCSI_ERR1	PERM	H	ADAPTER ERROR
069DB93B	MEM2	PERM	H	Memory failure
06ABB2EB	COM_CFG_BUST	PERM	S	Configuration failed: bad bus type
06CC7029	CXMA_CFG_FEPOS	PERM	S	Adapter FEPOS Execution Failed
0733FFA0	SDA_ERR2	TEMP	H	STORAGE SUBSYSTEM FAILURE
0734DA1D	DISKETTE_ERR3	PERM	H	DISKETTE MEDIA ERROR
08502E29	FDDI_TRACE	PERM	H	ADAPTER ERROR
0873CF9F	TTY_TTYHOG	TEMP	S	ttyhog over-run
087468D0	PSLA002	TEMP	S	SOFTWARE PROGRAM ERROR
08784A20	TOK_RMV_ADAP2	TEMP	S	REMOVE ADAPTER COMMAND RECEIVED
0A667C32	WHP0001	TEMP	S	SOFTWARE PROGRAM ERROR
0A940597	NB9	TEMP	S	SOFTWARE PROGRAM ERROR
0C1EC9FA	LVM_SA_WRTERR	UNKN	H	Failed to write Volume Group Status Area
0CACEC26	RS_PROG_IOCC	UNKN	S	Software error: iocc not configured
0CFAD921	RS_PROG_SLIH	UNKN	S	Software error: cannot find slih
0D5C1698	X25_ALERT33	PERM	H	X-33 (DCE) RESET INDICATION X.25 ADAPTER
0E017ED1	MEMORY	PERM	H	Memory failure
0E37FE58	LVM_BBEPOOL	UNKN	H	Bad block relocation failure - PV no lon
0EC7E7E5	EPOW_RES	UNKN	H	Electrical power resumed
0F27AAE5	CORE_DUMP	PERM	S	SOFTWARE PROGRAM ABNORMALLY TERMINATED
0F568474	IENT_ERR2	TEMP	S	CONFIGURATION OR CUSTOMIZATION ERROR
103F1912	X25_ALERT34	TEMP	H	X-34 (DCE) RESTART INDICATION X.25 ADAPT
10C6CED6	MPQP_RCVERR	TEMP	H	COMMUNICATION PROTOCOL ERROR
1104AA28	SYS_RESET	TEMP	S	System reset interrupt received
1251B5B7	LION_HRDWRE	PERM	S	Cannot access memory: 64 port controller
13881423	SCSI_ERR4	TEMP	H	MICROCODE PROGRAM ERROR
13C8A0AA	NB22	PERM	S	SOFTWARE PROGRAM ERROR
150ACBA4	X25_ALERT39	TEMP	H	X-39 (DCE) TIMEOUT ON CLEAR IND, T13
1581762B	DISK_ERR4	TEMP	H	DISK OPERATION ERROR
1588DDD9	CDROM_ERR3	PERM	H	OPTICAL DISK DRIVE ERROR
160544E1	SLA_DRIVER_ERR	PERM	H	SLA LINK CHECK fault in laser driver
1642B5A7	X25_ALERT26	PERM	H	X-26 TIMEOUT ON RESTART REQUEST, T20
173D5818	CDROM_ERR7	TEMP	H	OPTICAL DISK DRIVE ERROR
17A1F1E4	ACPA_INITZ	TEMP	S	Host independent initialization failed
18A546CD	LVM_BBDIRERR	UNKN	H	Bad block relocation failure - PV no lon
18B25E18	ACPA_INTR2	TEMP	S	Unexpected interrupt
192AC071	ERRLOG_OFF	TEMP	O	Error logging turned off
1A1D42F9	ACPA_LOAD	TEMP	S	Failed loading microcode
1A2E7186	LVM_MISSPVADDED	UNKN	S	Physical volume defined as missing
1A660730	C327_START	PERM	S	C327 Start error
1A9465A3	LVM_MWCWFAIL	UNKN	H	Mirror Write Cache write failed

```
1AC82784 LVM_SA_FRESHPP      UNKN S  Physical partition marked active
1B1647DF MPQP_XMTUND         PERF H  COMMUNICATIONS UNDERRUN
1CCD189F NB21                PERM S  SOFTWARE PROGRAM ERROR
1D5588BE WHP0013             TEMP S  SOFTWARE PROGRAM ERROR
1E629BB1 RS_8_16_ARB         PERM S  Invalid 8/16 port arbitration register
1F05D2DE FDDI_DWNLD          TEMP H  MICROCODE PROGRAM ABNORMALLY TERMINATED
1FD6C71A X25_ALERT32         PERM H  X-32 (DCE) CLEAR INDICATION X.25 ADAPTER
20188DE1 TOK_WIRE_FAULT      PERM H  WIRE FAULT
20FAED7F DSI_PROC            PERM S  Data Storage Interrupt, Processor
21D5B396 NB28                TEMP S  SOFTWARE PROGRAM ERROR
21F54B38 DISK_ERR1           PERM H  DISK OPERATION ERROR
225E3B63 KERNEL_PANIC        TEMP S  SOFTWARE PROGRAM ABNORMALLY TERMINATED
22F7B47B RS_MEM_IOCC         PERM S  Cannot allocate memory: iocc structure
233E36D2 NB26                PERM S  SOFTWARE PROGRAM ERROR
24247FB2 WHP0006             TEMP S  SOFTWARE PROGRAM ERROR
24DCDBA8 NB24                TEMP S  SOFTWARE PROGRAM ERROR
25D74748 EU_DIAG_ACC         PERM S  Cannot perform destructive diagnostics
270CB959 VCA_INTR2           TEMP S  Unexpected interrupt
273FE0AC NB14                PERM S  SOFTWARE PROGRAM ERROR
27C1EFFF DSI_IOCC            PERM H  Data Storage Interrupt, IOCC
28935927 NLS_MAP             PERM S  SOFTWARE PROGRAM ERROR
289590AE NB13                PERM S  SOFTWARE PROGRAM ERROR
2920ZCA2 COM_MEM_SLIH        PERM S  Cannot allocate memory: slih structure
2929FD6D FDDI_RCVRY_EXIT     TEMP H  PROBLEM RESOLVED
29975223 COM_CFG_DEVD        PERM S  Configuration failed: devswdel failed
2A53071F FDDI_PATH_ERR       PERM H  ADAPTER ERROR
2A7392A2 COM_CFG_MNR         PERM S  Configuration failed: bad minor number
2AA90CCD CXMA_IO_ATT         PERM S  I/O Segment Attach Failed
2B60DD24 WHP0012             TEMP S  SOFTWARE PROGRAM ERROR
2B76062D MPQP_BFR            PERF S  OUT OF RESOURCES
2BFA76F6 REBOOT_ID           TEMP S  System shutdown by user
2C7CE30E EU_BAD_ADPT         PERM H  Expansion unit error
2CF9AB6C CFGMGR_MEMORY       UNKN S  Not enough memory for configuration mgr
2D3BDDD6 BADISK_ERR8         PERM H  DISK OPERATION ERROR
2DACEE65 FDDI_ADAP_CHECK     PERM H  ADAPTER ERROR
2F24221A ENT_ERR4            TEMP H  ADAPTER ERROR
2F65D788 X25_ALERT7          PERM H  X-7  MODEM FAILURE: ACU NOT RESPONDING
30911E21 X25_ALERT5          PERM H  X-5  MODEM FAILURE: DCD, DSR, CABLE
30F182A4 CDROM_ERR1          PERM H  OPTICAL DISK OPERATION ERROR
342CB115 FDDI_TX_ERR         TEMP H  ADAPTER ERROR
345707F5 TTY_INTR_HOG        TEMP H  PIO exception
34FC3203 CDROM_ERR2          TEMP H  OPTICAL DISK OPERATION ERROR
3503BDBA X25_ALERT30         PERM H  X-30 DIAGNOSTIC PACKET RECEIVED
35890E9F TOK_NOMBUFS         UNKN S  RESOURCE UNAVAILABLE
358D0A3E DOUBLE_PANIC        TEMP S  SOFTWARE PROGRAM ABNORMALLY TERMINATED
35BE4BC0 IENT_ERR1           TEMP H  ADAPTER ERROR
35BFC499 DISK_ERR3           PERM H  DISK OPERATION ERROR
36C3328B ATE_ERR1            PERM S  COMMUNICATION PROTOCOL ERROR
3766B2C7 FDDI_BYPASS         PERM H  ADAPTER ERROR
384E0485 BADISK_ERR1         TEMP H  DISK OPERATION ERROR
39DCD110 SLA_PROG_ERR        TEMP S  SLA programming check
3A30359F INIT_RAPID          TEMP S  SOFTWARE PROGRAM ERROR
3A58ABE2 RS_PIN_IOCC         PERM S  Cannot pin memory: iocc structure
3A67AFE0 ATE_ERR6            PERM S  COMMUNICATION PROTOCOL ERROR
3A9C2352 DISKETTE_ERR2       UNKN H  DISKETTE DEVICE FAILURE
3B145117 IENT_ERR4           UNKN S  UNDETERMINED ERROR
3C19F251 NB2                 TEMP S  SOFTWARE PROGRAM ERROR
3CFF4028 DISK_ERR5           UNKN H  UNDETERMINED ERROR
3D858A1B MEM1                PERM H  Memory failure
```

```
3EC3C657 COM_CFG_NADP       PERM S  Configuration failed: adapter missing
3F86401A LION_BOX_DIED      PERM H  Lost communication: 64 port concentrator
419D40C2 NB23               PERM S  SOFTWARE PROGRAM ERROR
4224BA8C WHP0008            TEMP S  SOFTWARE PROGRAM ERROR
4287A984 COM_CFG_BUSID      PERM S  Configuration failed: bad bus id range
43D4ADCE TTY_PARERR         TEMP S  Parity/Framing error on input
44CB9ECE MPQP_DSRTO         TEMP H  UNABLE TO COMMUNICATE WITH DEVICE
4523CAA9 CMDLVM             PERF H  DISK OPERATION ERROR
476B351D TAPE_ERR2          PERM H  TAPE DRIVE FAILURE
47E84916 IENT_ERR5          UNKN S  COMMUNICATIONS SUBSYSTEM FAILURE
484F5514 NB6                TEMP S  SOFTWARE PROGRAM ERROR
4865FA9B TAPE_ERR1          PERM H  TAPE OPERATION ERROR
4A29D32A MACHINECHECK       PERM H  Machine Check
4A4FBE2B NB16               TEMP S  SOFTWARE PROGRAM ERROR
4AB56573 CAT_ERR2           PERM S  MICROCODE PROGRAM ERROR
4B0E39BB CXMA_MEM_CH        PERM S  Can¢t Allocate ch_t Structures
4C2BDA1E NB3                TEMP S  SOFTWARE PROGRAM ERROR
4CEBE931 COM_CFG_UIO        PERM S  Configuration failed: resid not correct
4EDEF5A1 SCSI_ERR5          PERM S  SOFTWARE PROGRAM ERROR
4F3E9630 INIT_UNKNOWN       TEMP S  SOFTWARE PROGRAM ERROR
4F515DF0 WHP0005            TEMP S  SOFTWARE PROGRAM ERROR
504B04D3 NB18               PERM S  SOFTWARE PROGRAM ERROR
506E5213 ACPA_IOCTL2        TEMP S  Invalid ioctl request
50CA5315 LION_BUFFER0       TEMP S  Buffer overrun: 64 port concentrator
5114C792 COM_CFG_IFLG       PERM S  Configuration failed: bad interrupt flag
51F9313A NB17               PERM S  SOFTWARE PROGRAM ERROR
52DB7218 SCSI_ERR6          TEMP S  SOFTWARE PROGRAM ERROR
532D1C49 TOK_DOWNLOAD       PERM H  MICROCODE PROGRAM ABNORMALLY TERMINATED
53920B1F ACPA_IOCTL1        PERM S  Invalid ioctl request
5416CE51 COM_TEMP_PIO       TEMP H  PIO exception
544FF289 COM_CFG_SLIH       PERM S  Configuration failed: i_init of slih
54B73180 LVM_BBDIRFUL       UNKN H  Bad block relocation failure
54E423ED SCSI_ERR9          PERM H  Potential data loss condition
5529E45B X25_ALERT21        PERM H  X-21 CLEAR INDICATION RECEIVED
5537AC5F TAPE_ERR4          PERM H  TAPE DRIVE FAILURE
56816728 MPQP_CTSTO         TEMP H  COMMUNICATION PROTOCOL ERROR
57797644 X25_ADAPT          PERM H  ADAPTER ERROR
592D5E9D TOK_WRAP_TST       PERM H  OPEN FAILURE
59792439 X25_ALERT12        TEMP H  X-12 FRAME TYPE Z RECEIVED
59853D4A CXMA_CFG_TALLOC    PERM S  talloc failed
59D54E37 X25_ALERT16        TEMP H  X-16 FRAME TYPE Z SENT
5A48B4FF FDDI_RCVRY_TERM    PERM H  ADAPTER ERROR
5AE97EAA MSLA_PROTOCOL      TEMP S  COMMUNICATION PROTOCOL ERROR
5CC986A0 SCSI_ERR3          PERM H  MICROCODE PROGRAM ERROR
5CE03B80 INIT_OPEN          TEMP S  SOFTWARE PROGRAM ERROR
5CFBFA4A WHP0004            TEMP S  SOFTWARE PROGRAM ERROR
5D1F16FA CAT_ERR8           TEMP H  ADAPTER ERROR
5D66BBC4 DUMP_STATS         UNKN S  System dump
5DFEADCB LVM_HWREL          UNKN H  Hardware disk block relocation achieved
5E9573AA CXMA_ERR_ASSRT     PERM S  Driver Assert Message
5F504A40 SLA_SIG_ERR        PERM H  SLA LINK CHECK signal failure
60D5349F COM_PIN_SLIH       PERM S  Cannot pin memory: slih structure
618DB24A X25_ALERT24        PERM H  X-24 CLEAR REQUEST BY X.25 ADAPTER
627A4F55 BADISK_ERR3        TEMP H  DISK OPERATION ERROR
6297CA97 DUMP               TEMP H  Dump device error
66C3412B RS_MEM_EDGE        PERM S  Cannot allocate memory: edge structure
680A6C7C CXMA_CFG_PORT      PERM S  Bad Adapter I/O Port Address
684B0E5C LVM_BBDIR90        UNKN H  Bad block directory over 90% full
68F9701C CXMA_ADP_FAIL      PERM H  Async Adapter Failed
```

```
69221791 MSLA_START          TEMP S  OUT OF RESOURCES
6B0B47FA CFGMGR_LOCK         UNKN S  Could not acquire configuration lock
6D6B57F9 TOK_BAD_ASW         PERM H  MICROCODE PROGRAM ERROR
6F7D7290 X25_ALERT11         TEMP H  X-11 FRAME TYPE Y RECEIVED
6FD1189E X25_ALERT15         TEMP H  X-15 FRAME TYPE Y SENT
70559CAE NB4                 PERM S  SOFTWARE PROGRAM ERROR
71248BF5 ISI_PROC            PERM S  Instruction Storage Interrupt
7239AC3D FDDI_LLC_ENABLE     TEMP H  PROBLEM RESOLVED
72CBC436 TMSCSI_UNKN_SFW_ERR UNKN S  SOFTWARE PROGRAM ERROR
74533D1A EPOW_SUS            UNKN H  LOSS OF ELECTRICAL POWER
74E0CEA8 X25_IPL             PERM H  ADAPTER ERROR
760470A6 IENT_ERR3           TEMP H  Data Storage Interrupt, IOCC
76C9D063 DSI_SLA             PERM H  Data Storage Interrupt, SLA
770F9606 BADISK_ERR2         PERM H  DISK OPERATION ERROR
773D6C8E NB7                 TEMP S  SOFTWARE PROGRAM ERROR
77E0148A MEM3                PERM H  Memory failure
7873CE72 X25_ALERT31         PERM H  X-31 RESET INDICATION PACKET RECEIVED
794A4421 X25_ALERT37         TEMP H  X-37 (DCE) TIMEOUT ON RESET IND, T12
7993098B COM_CFG_UNK         PERM S  Configuration failed: bad adapter type
79FED1ED NB29                PERM S  SOFTWARE PROGRAM ERROR
7A9C71E6 X25_ALERT18         PERM H  X-18 UNEXPECTED DISC RECEIVED
7A9E20BB MPQP_XFTO           PERM H  ADAPTER ERROR
7AB881D9 MISC_ERR            UNKN H  Miscellaneous interrupt
7B3D4206 SLA_EXCEPT_ERR      PERM H  Internal serial link adapter exception
7BDD117A TOK_RCVRY_ENTER     TEMP H  ADAPTER ERROR
7C197591 SLA_FRAME_ERR       TEMP H  SLA LINK CHECK possible lost frame
7D1E4727 TOK_DUP_ADDR        TEMP S  OPEN FAILURE
7EF0A4FF CFGMGR_NONFATAL_DB  UNKN S  Configuration mgr nonfatal database err
7F0052C6 COM_CFG_UNPIN       PERM S  Configuration failed: unpincode failed
7FF45EC0 WHP0003             TEMP S  SOFTWARE PROGRAM ERROR
804055EB NB15                PERM S  SOFTWARE PROGRAM ERROR
804C1878 COM_CFG_RESID       PERM S  Configuration failed: resid not correct
80A357F9 INIT_CREATE         TEMP S  SOFTWARE PROGRAM ERROR
80F672FF CAT_ERR4            TEMP S  RESOURCE UNAVAILABLE
813E4B9A X25_ALERT10         TEMP H  X-10 FRAME TYPE X RECEIVED
81922194 X25_ALERT14         TEMP H  X-14 FRAME TYPE X SENT
835C5977 ACPA_INTR1          TEMP S  Interrupt handler registration failed
836A2443 X25_CONFIG          PERM H  X.25 CONFIGURATION ERROR
83E4C0B2 LVM_SWREL           UNKN H  Software disk block relocation achieved
84917289 LVM_BBRELMAX        UNKN H  Bad block relocation failure - PV no lon
84EE0148 MPQP_QUE            TEMP H  MPQP unable to access queue
861365E7 EU_CFG_NPLN         PERM S  Configuration failed: adapter missing
868921F2 TMSCSI_READ_ERR     TEMP H  Attached SCSI initiator error
86922CCD X25_ALERT27         PERM H  X-27 TIMEOUT ON RESET REQUEST, T22
89B52AA5 CONSOLE             PERM S  SOFTWARE PROGRAM ERROR
89C695BB ACPA_INTR4          TEMP S  Interrupt timed out
8B5D61E6 CXMA_MEM_TTY        PERM S  Can¢t Allocate tty_t Structures
8BBE428E TOK_BEACON3         TEMP S  TOKEN-RING TEMPORARY ERROR
8C0353CB MPQP_X21CECLR       PERM S  X.21 ERROR
8D2CC3AA MSLA_WRITE          TEMP S  ADAPTER ERROR
8DCE65AF FDDI_MC_ERR         TEMP H  ADAPTER ERROR
8DD34341 CDROM_ERR6          TEMP H  OPTICAL DISK DRIVE ERROR
8EA094FF CHECKSTOP           TEMP H  Checkstop
8FEF9795 DISKETTE_ERR6       PERM H  PIO exception
904C6053 VCA_INTR1           TEMP S  Interrupt handler registration failed
9060A2F8 CAT_ERR3            TEMP S  RESOURCE UNAVAILABLE
90809FD9 TOK_ERR10           PERM H  MANAGEMENT SERVER REPORTING LINK ERROR
91D6C4F8 CDROM_ERR5          PERM H  OPTICAL DISK DRIVE ERROR
91E8D590 TOK_MC_ERR          PERM H  ADAPTER ERROR
```

```
91F9700D LVM_SA_QUORCLOSE    UNKN H  Quorum lost, volume group closing
91FDA5E4 CFGMGR_OPTION       UNKN S  Invalid option: configuration manager
925A4C9B SLA_CRC_ERR         TEMP H  SLA LINK CHECK crc error
92A72C14 COM_CFG_ILVL        PERM S  Configuration failed: interrupt level
9359F226 LVM_MISSPVRET       UNKN S  Physical volume is now active
974CC901 X25_ALERT19         PERM H  X-19 DM RXD DURING LINK ACTIVATION
9844042C NB27                TEMP S  SOFTWARE PROGRAM ERROR
98A70F55 ENT_ERR5            UNKN S  RESOURCE UNAVAILABLE
98F39A90 TMSCSI_RECVRD_ERR   TEMP H  Attached SCSI target device error
99227331 ENT_ERR3            PERM H  ADAPTER ERROR
9A335282 EXCHECK_RSC         PERM H  External Check, DMA
9AD6AC9F VCA_INTR4           TEMP S  Interrupt timed out
9B55A553 FDDI_RMV_ADAP       PERM H  REMOVE ADAPTER COMMAND RECEIVED
9C7FE90B LION_MEM_ADAP       PERM S  Cannot allocate memory: adap structure
9D30B78E TTY_OVERRUN         TEMP S  Receiver over-run on input
9DBCFDEE ERRLOG_ON           TEMP O  Error logging turned on
9E45396D NB5                 TEMP S  SOFTWARE PROGRAM ERROR
A194D797 TOK_ERR15           UNKN H  ADAPTER ERROR
A28B68BD MSLA_ADAPTER        PERM H  ADAPTER ERROR
A386E435 ENT_ERR1            PERM H  ADAPTER ERROR
A38E8CF2 CDROM_ERR4          TEMP H  OPTICAL DISK DRIVE ERROR
A5417864 WHP0011             TEMP S  SOFTWARE PROGRAM ERROR
A668F553 DISK_ERR2           PERM H  DISK OPERATION ERROR
A6BAD8E6 CORRECTED_SCRUB     TEMP H  Memory scrubbing corrected ECC error
A741AD52 MPQP_DSROFFTO       TEMP H  UNABLE TO COMMUNICATE WITH DEVICE
A80659F3 WHP0014             TEMP S  SOFTWARE PROGRAM ERROR
A84C681B VCA_MEM             TEMP S  Failed pinning memory
A853F9CE EU_DIAG_MEM         PERM S  Cannot allocate memory: wrap buffer
A92AE715 DISKETTE_ERR1       TEMP H  DISKETTE OPERATION ERROR
A9844FEE EXCHECK_DMA         PERM H  External Check, DMA
A9ED5BB6 SDC_ERR1            PERM H  LINK ERROR
AA8AB241 OPMSG               TEMP O  OPERATOR NOTIFICATION
AAD5C121 TOK_AUTO_RMV        PERM H  AUTO REMOVAL
ABB81CD5 ENT_ERR2            TEMP H  COMMUNICATION PROTOCOL ERROR
ABEC9F35 TOK_RMV_ADAP1       PERM H  OPEN FAILURE
AC47FA8A X25_ALERT38         TEMP H  X-38 (DCE) TIMEOUT ON CALL IND, T11
ACDAE3FC TOK_ADAP_CHK        PERM H  UNABLE TO COMMUNICATE WITH DEVICE
AD682624 CDROM_ERR8          UNKN H  UNDETERMINED ERROR
AD917FBA MPQP_ASWCHK         PERM S  MICROCODE PROGRAM ERROR
AEC7B1B0 TOK_BEACON2         PERM H  TOKEN-RING INOPERATIVE
AFF4BD94 NB30                PERM S  SOFTWARE PROGRAM ERROR
B135AE8B SDA_ERR1            PERM H  STORAGE SUBSYSTEM FAILURE
B1462F15 SDC_ERR3            TEMP H  STORAGE SUBSYSTEM FAILURE
B18287F3 SDA_ERR4            TEMP H  UNDETERMINED ERROR
B188909A LVM_SA_STALEPP      UNKN S  Physical partition marked stale
B216DB3E COM_CFG_PORT        PERM S  Configuration failed: port configured
B29547EF CXMA_CFG_RST        PERM S  Adapter Reset Failed
B3683B72 FDDI_XCARD          PERM H  ADAPTER ERROR
B5982183 EU_CFG_BUSY         PERM S  Configuration failed: in use
B598ECB3 PSLA001             TEMP H  DEVICE ERROR
B617E928 TAPE_ERR6           TEMP H  TAPE OPERATION ERROR
B63E9C5E RS_BAD_INTER        PERM S  Interrupt from non-existant port
B6A6F2B7 CXMA_CFG_MPORT      PERM S  Bad or Missing Port on Adapter
B7164FA8 WHP0007             TEMP S  SOFTWARE PROGRAM ERROR
B73A1D33 X25_ALERT13         TEMP H  X-13 FRAME TYPE W RECEIVED
B73BC3CD DISKETTE_ERR4       UNKN H  DISKETTE OPERATION ERROR
B76A0A99 LION_CHUNKNUMC      TEMP S  Bad chunk count: 64 port controller
B7BF9C85 CXMA_CFG_MEM        PERM S  Bad Adapter Memory Address
B7F0EC53 NB10                TEMP S  SOFTWARE PROGRAM ERROR
```

```
B8892A14 DSI_SCU             PERM H  Data Storage Interrupt, SCU
BAB1383B NB8                 TEMP S  SOFTWARE PROGRAM ERROR
BAECC981 SDM_ERR1            PERM H  MICROCODE PROGRAM ERROR
BB5C513F ACPA_MEM            TEMP S  Failed pinning memory
BBA1D78B ACPA_UCODE          TEMP S  Failed loading microcode onto M-ACPA/A
BC8F0BBB COM_CFG_DEVA        PERM S  Configuration failed: devswadd failed
BDA444C8 SLA_PARITY_ERR      TEMP H  SLA buffer parity error
BE42630E REPLACED_FRU        PERM H  Repair action
BE7E5290 LION_PIN_ADAP       PERM S  Cannot pin memory: adap structure
BE7F0C5D COM_CFG_DMA         PERM S  Configuration failed: dma level conflict
BE910C7F CAT_ERR7            TEMP S  RESOURCE UNAVAILABLE
BF06FA0D FDDI_LLC_DISABLE    TEMP H  LAN ERROR
BF3F8438 PSLA003             TEMP H  LINK ERROR
BF6D9219 LION_UNKCHUNK       TEMP S  Unknown error code: 64 port concentrator
BF93B600 TOK_RCVRY_TERM      PERM H  ADAPTER ERROR
BFEA74DC CXMA_MEM_ATT        PERM S  Memory Segment Attach Failed
C0073BB4 TTY_BADINPUT        TEMP S  Bad ttyinput return
C0514A3F X25_ALERT35         TEMP H  X-35 (DCE) RESTART RESET RECEIVED
C1423E5B WHP0010             TEMP S  SOFTWARE PROGRAM ERROR
C14C511C SCSI_ERR2           TEMP H  ADAPTER ERROR
C2B80BFB X25_ALERT36         TEMP H  X-36 (DCE) TIMEOUT ON RESTART IND, T10
C580DED6 WHP0009             TEMP S  SOFTWARE PROGRAM ERROR
C5C09FFA PGSP_KILL           PERM S  SOFTWARE PROGRAM ABNORMALLY TERMINATED
C67E7D0F LVM_HWFAIL          UNKN H  Hardware disk block relocation failed
C6ACA566 SYSLOG              UNKN S  Message redirected from syslog
C6EB3E75 FDDI_SELF_TEST      TEMP H  LAN ERROR
C70E1E46 X25_ALERT17         PERM H  X-17 FRAME RETRY N2 REACHED
C88D3DD8 MPQP_X21CPS         PERM S  X.21 ERROR
C89DE914 C327_INTR           PERM S  C327 Interrupt error
C8F22E8E FLPT_UNAVAIL        PERM S  OPERATOR NOTIFICATION
C92F456F NB11                TEMP S  SOFTWARE PROGRAM ERROR
C9A0C741 X25_UCODE           PERM H  X.25 MICROCODE ERROR
C9E358D3 CXMA_LINE_ERR       PERM H  Synchronous Line Errors
C9F4EE17 EU_CFG_NADP         PERM S  Configuration failed: adapter missing
CBE1D1A5 LVM_SA_PVMISS       UNKN H  Physical volume declared missing
CBE25456 MSLA_INTR           TEMP S  COMMUNICATION PROTOCOL ERROR
CEDCB90F FDDI_PIO            TEMP H  PIO exception
CF4781D3 BADISK_ERR4         PERM H  DISK OPERATION ERROR
CFC1A4DD MPQP_ADPERR         PERM H  ADAPTER ERROR
CFCDE8F6 FDDI_DOWN           TEMP H  ADAPTER ERROR
CFFF77BD TOK_ADAP_ERR        PERM H  Potential data loss condition
D080E08D CAT_ERR5            TEMP H  ADAPTER ERROR
D2360951 TOK_CONGEST         PERF S  COMMUNICATIONS OVERRUN
D2B9B5A9 BADISK_ERR5         PERM H  DISK OPERATION ERROR
D3B0ECBF X25_ALERT8          PERM H  X-8  X.21 NOT CONNECTED
D3F26EC3 NB1                 TEMP S  SOFTWARE PROGRAM ERROR
D41B92E8 RS_PIN_EDGEV        PERM S  Cannot pin memory: edge vector
D62AAFD8 LVM_BBDIRBAD        UNKN H  Bad block relocation failure - PV no lon
D7BDE2AD INTR_ERR            UNKN H  UNDETERMINED ERROR
D7DDDC46 CAT_ERR1            PERM H  MICROCODE PROGRAM ABNORMALLY TERMINATED
D824DB48 VCA_INTR3           TEMP S  Invalid interrupt
D84B1C5B LION_MEM_LIST       PERM S  Cannot allocate memory: ttyp_t list
D8EA614B FDDI_USYS           UNKN S  UNDETERMINED ERROR
D9EE4AC1 EU_CFG_GONE         PERM S  Configuration failed: unconfigured
DA244DCA COM_CFG_PIN         PERM S  Configuration failed: pincode failed
DA80B2D4 NB12                PERM S  SOFTWARE PROGRAM ERROR
DB3E3DFD ENT_ERR6            PERM H  CSMA/CD LAN COMMUNICATIONS LOST
DB451F82 MPQP_RCVOVR         PERF H  COMMUNICATIONS OVERRUN
DBF56911 EU_CFG_HERE         PERM S  Configuration failed: already configured
```

```
DBF832FF LVM_BBFAIL          UNKN H  Bad block relocation failure - PV no lon
DD0E4902 TOK_RCVRY_EXIT      TEMP H  PROBLEM RESOLVED
DD11B4AF PROGRAM_INT         PERM S  Program Interrupt
DD2201A9 X25_ALERT28         PERM H  X-28 TIMEOUT ON CALL REQUEST, T21
DDBCA0EE VCA_IOCTL2          TEMP S  Invalid ioctl request
DFC508F5 PPRINTER_ERR1       UNKN H  PRINTER ERROR
E0EA14BF TOK_BEACON1         TEMP S  OPEN FAILURE
E180FD0E CXMA_CONC_DOWN      PERM H  Concentrator Removed From System
E18E984F SRC                 PERM S  SOFTWARE PROGRAM ERROR
E2109F7A COM_PERM_PIO        PERM H  PIO exception
E225351D CXMA_ERR_EVNT       PERM S  Event handler Failure
E252FE92 MPQP_X21DTCLR       PERM S  X.21 ERROR
E2A4EC26 RS_MEM_EDGEV        PERM S  Cannot allocate memory: edge vector
E2B9E02B TTY_PROG_PTR        UNKN S  Software error: t_hptr field invalid
E47E212E INIT_UTMP           TEMP S  SOFTWARE PROGRAM ERROR
E4EF0A90 WHP0002             TEMP S  SOFTWARE PROGRAM ERROR
E4F5F86E MPQP_IPLTO          PERM H  ADAPTER ERROR
E61501A6 MPQP_X21TO          TEMP H  X.21 ERROR
E64EC259 TAPE_ERR3           PERM H  TAPE DRIVE FAILURE
E6599C95 X25_ALERT23         PERM H  X-23 RESET REQUEST BY X.25 ADAPTER
E6784BC4 X25_ALERT29         PERM H  X-29 TIMEOUT ON CLEAR REQUEST, T23
E6CDBCFC CFGMGR_PROGRAM_NF   UNKN S  Program or method not found
E70473E7 VCA_IOCTL1          PERM S  Invalid ioctl request
E79A3C09 ACPA_INTR3          TEMP S  Invalid interrupt
E7D0FE3F RS_PIN_EDGE         PERM S  Cannot pin memory: edge structure
E7E2E3E9 NLS_BADMAP          PERM S  Software error:  NLS map corrupted
E85C5C4C HFTERR              PERM S  SOFTWARE PROGRAM ERROR
E9645CC5 FDDI_RCV            UNKN H  ADAPTER ERROR
E97374FF RS_MEM_PVT          PERM S  Cannot allocate memory: priv. structure
EA388E60 X25_ALERT22         PERM H  X-22 RESTART INDICATION RECEIVED
EB5F98B2 RCMERR              PERM S  SOFTWARE PROGRAM ERROR
EE18DF01 TMSCSI_CMD_ERR      TEMP H  Attached SCSI target device error
EE8BC5D8 CXMA_CFG_BIOS       PERM S  Adapter BIOS Initialization Failed
EFEC314D DISKETTE_ERR5       TEMP H  PIO exception
F15F3C50 FDDI_RCVRY_ENTER    PEND H  Recovery logic initiated by device
F2F30ADF FDDI_PORT           TEMP H  ADAPTER ERROR
F3D17657 CXMA_CFG_MTST       PERM S  Adapter Memory Test Failed
F438E969 SDC_ERR2            PERM H  STORAGE SUBSYSTEM FAILURE
F4CB727F FDDI_SELFT_ERR      PERM H  ADAPTER ERROR
F5345AAB NB25                PERM S  SOFTWARE PROGRAM ERROR
F5458763 COM_CFG_ADPT        PERM S  Configuration failed: already configured
F6E3C547 ATE_ERR7            TEMP S  COMMUNICATION PROTOCOL ERROR
F734B194 NB19                PERM S  SOFTWARE PROGRAM ERROR
F7E70B81 EXCHECK_SCRUB       PERM H  OPERATOR NOTIFICATION
F81946D8 CFGMGR_CHILD        UNKN S  Configuration mgr child process failed
F9171B5C CFGMGR_FATAL_DB     UNKN S  Configuration mgr fatal database problem
F924E95E TOK_PIO_ERR         PERM H  ADAPTER ERROR
FB683A72 ACCT_OFF            TEMP S  EC26
FBD2B2B5 MSLA_IOCTL          TEMP S  ADAPTER ERROR
FBF0BFC1 TMSCSI_UNRECVRD_ERR PERM H  Attached SCSI target device error
FCA960CE TOK_ESERR           TEMP S  EXCESSIVE TOKEN-RING ERRORS
FDE6A5A1 COM_CFG_BUSI        PERM S  Configuration failed: bad bus ID
FE1DA20A TOK_ERR5            PERM H  OPEN FAILURE
FE6A2D60 COM_CFG_INTR        PERM S  Configuration failed: interrupt priority
FEC31570 SDA_ERR3            PERM H  UNDETERMINED ERROR
FED1497C MSLA_CLOSE          TEMP S  SOFTWARE PROGRAM ABNORMALLY TERMINATED
FFC9ECAA TOK_TX_ERR          PERM H  ADAPTER ERROR
FFE2F73A TAPE_ERR5           UNKN H  UNDETERMINED ERROR
```

Appendix D. Disk Setup in an HACMP Cluster

This appendix gives detailed descriptions of the setup of different kinds of shared disk devices for HACMP. You will see how cluster nodes are connected to shared disks and how the storage space on these devices becomes visible to the operating system.

The appendix is divided into three sections, each of which deals with a particular type of disk or subsystem. These sections are:

- SCSI disks and subsystems
- RAID subsystems
- 9333 Serial disk subsystems
- Serial Storage Architecture (SSA) disk subsystems

D.1 SCSI Disks and Subsystems

The SCSI adapters that can be used on a shared SCSI bus in an HACMP cluster are:

- SCSI-2 Differential Controller (FC: 2420, PN: 43G0176)
- SCSI-2 Differential Fast/Wide Adapter/A (FC: 2416, PN: 65G7315)
- Enhanced SCSI-2 Differential Fast/Wide Adapter/A (FC: 2412, PN: 52G3380)

  (This adapter was only supported under AIX 4.1 and HACMP 4.1 for AIX at the time of publishing, but testing was underway to certify the adapter under HACMP/6000 Version 3.1)

The non-RAID SCSI disks and subsystems that you can connect as shared disks in an HACMP cluster are:

- 7204 Models 215, 315, 317, and 325 External Disk Drives
- 9334 Models 011 and 501 SCSI Expansion Units
- 7134-010 High Density SCSI Disk Subsystem

D.1.1 SCSI Adapters

The SCSI-2 Differential Controller is used to connect to 8-bit disk devices on a shared bus. The SCSI-2 Differential Fast/Wide Adapter/A or Enhanced SCSI-2 Differential Fast/Wide Adapter/A is usually used to connect to 16-bit devices but can also be used with 8-bit devices.

In a dual head-of-chain configuration of shared disks, there should be no termination anywhere on the bus except at the extremities. Therefore, you should remove the termination resistor blocks from the SCSI-2 Differential Controller and the SCSI-2 Differential Fast/Wide Adapter/A or Enhanced SCSI-2 Differential Fast/Wide Adapter/A. The positions of these blocks (U8 and U26 on the SCSI-2 Differential Controller, and RN1, RN2 and RN3 on the SCSI-2 Differential Fast/Wide Adapter/A and Enhanced SCSI-2 Differential Fast/Wide Adapter/A) are shown in Figure 5 on page 108 and Figure 6 on page 108 respectively.

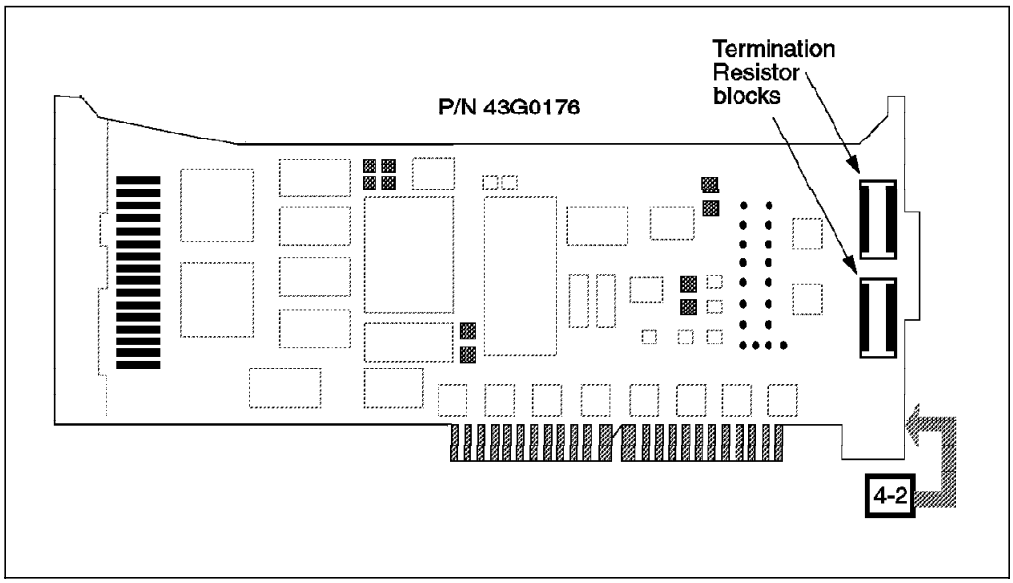

Figure 5. Termination Resistor Blocks on the SCSI-2 Differential Controller

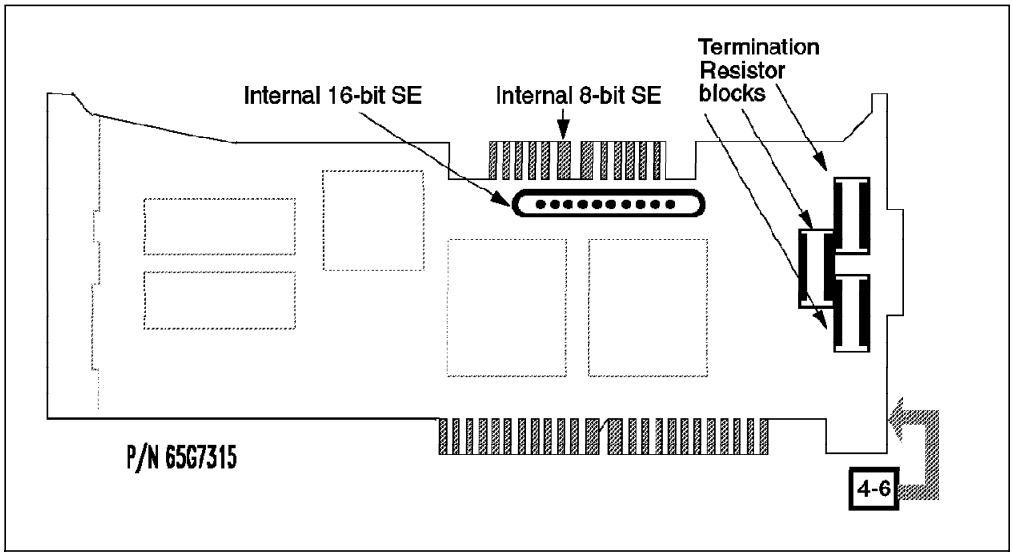

Figure 6. Termination Resistor Blocks on the SCSI-2 Differential Fast/Wide Adapter/A and Enhanced SCSI-2 Differential Fast/Wide Adapter/A

The ID of a SCSI adapter, by default, is 7. Since each device on a SCSI bus must have a unique ID, the ID of at least one of the adapters on a shared SCSI bus has to be changed.

The procedure to change the ID of a SCSI-2 Differential Controller is:

1. At the command prompt, enter `smit chgscsi`.
2. Select the adapter whose ID you want to change from the list presented to you.

```
                                SCSI Adapter

  Move cursor to desired item and press Enter.

    scsi0 Available 00-02 SCSI I/O Controller
    scsi1 Available 00-06 SCSI I/O Controller
    scsi2 Available 00-08 SCSI I/O Controller
    scsi3 Available 00-07 SCSI I/O Controller

  F1=Help                 F2=Refresh              F3=Cancel
  F8=Image                F10=Exit                Enter=Do
  /=Find                  n=Find Next
```

3. Enter the new ID (any integer from 0 to 7) for this adapter in the Adapter card SCSI ID field. Since the device with the highest SCSI ID on a bus gets control of the bus, set the adapter's ID to the highest available ID. Set the Apply change to DATABASE only field to **yes**.

```
                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                       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                      yes                   +



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

4. Reboot the machine to bring the change into effect.

The same task can be executed from the command line by entering:

```
# chdev -l scsi1 -a id=6 -P
```

Also with this method, a reboot is required to bring the change into effect.

The procedure to change the ID of a SCSI-2 Differential Fast/Wide Adapter/A or Enhanced SCSI-2 Differential Fast/Wide Adapter/A is almost the same as the one described above. Here, the adapter that you choose from the list you get after executing the smit chgsys command should be an ascsi device. Also, as, shown below, you need to change the external SCSI ID only.

```
                Change/Show Characteristics of a SCSI Adapter


  SCSI adapter                                          ascsi1
  Description                                           Wide SCSI I/O Control>
  Status                                                Available
  Location                                              00-06
  Internal SCSI ID                                      7                       +#
  External SCSI ID                                      [6]                     +#
  WIDE bus enabled                                      yes                     +
  ...
  Apply change to DATABASE only                         yes
```

The command line version of this is:

```
# chdev -l ascsi1 -a id=6 -P
```

As in the case of the SCSI-2 Differential Controller, a system reboot is required to bring the change into effect.

The maximum length of the bus, including any internal cabling in disk subsystems, is limited to 19 meters for buses connected to the SCSI-2 Differential Controller, and to 25 meters for those connected to the SCSI-2 Differential Fast/Wide Adapter/A or Enhanced SCSI-2 Differential Fast/Wide Adapter/A.

D.1.2 Individual Disks and Enclosures

The 7204-215 External Disk Drive is an 8-bit disk that can be connected to the SCSI-2 Differential Controller, the SCSI-2 Differential Fast/Wide Adapter/A, or the Enhanced SCSI-2 Differential Fast/Wide Adapter/A. While there is a theoretical limit of six such disks in an I/O bus connected to two nodes, HACMP supports up to four in a single bus. This support limit is based only on what has been specifically tested by development.

As there are typically choices to be made in lengths of cable connecting disks and adapters in the bus, it is important to keep in mind the bus length limits stated in the last section, while configuring your hardware.

The 7204 Model 315, 317, and 325 External Disk Drives are 16-bit disks that can only be connected to the SCSI-2 Differential Fast/Wide Adapter/A or Enhanced SCSI-2 Differential Fast/Wide Adapter/A. For HACMP, the tested limit of these disks in a single shared 16-bit bus is six for the 7204-315, and fourteen for the 7204-317 and 7204-325.

The 9334 Model 011 and 501 SCSI Expansion Units can each contain up to four 8-bit disks. Because of the bus length limitation, you can daisy-chain a maximum of two such units on a shared bus. The number of disks in the enclosures is determined by the number of free SCSI IDs in the bus. The enclosure itself does not have any SCSI ID.

The 7134-010 High Density SCSI Disk Subsystem can contain up to six 16-bit disks in the base unit and six more in the expansion unit. You can either configure your 7134 with just the base unit connected to one shared SCSI bus, or you can configure it with the base and the expansion unit attached to two different shared SCSI buses. The maximum number of disks in each unit is

determined by the number of available SCSI IDs on the shared bus to which it is attached.

D.1.3 Hooking It All Up

In this section we will list the different components required to connect SCSI disks and enclosures on a shared bus. We will also show you how to connect these components together.

D.1.3.1 7204-215 External Disk Drive

To connect a set of 7204-215s to SCSI-2 Differential Controllers on a shared SCSI bus, you need the following:

- SCSI-2 Differential Y-Cable

  FC: 2422 (0.765m), PN: 52G7348

- SCSI-2 Differential System-to-System Cable

  FC: 2423 (2.5m), PN: 52G7349

  This cable is used only if there are more than two nodes attached to the same shared bus.

- SCSI-2 DE Controller Cable

  FC: 2854 or 9138 (0.6m), PN: 87G1358 - OR -

  FC: 2921 or 9221 (4.75m), PN: 67G0593

- SCSI-2 DE Device-to-Device Cable

  FC: 2848 or 9134 (0.66m), PN: 74G8511

- Terminator

  Included in FC 2422 (Y-Cable), PN: 52G7350

Figure 7 shows four RS/6000s, each represented by one SCSI-2 Differential Controller, connected on an 8-bit bus to a chain of 7204-215s.

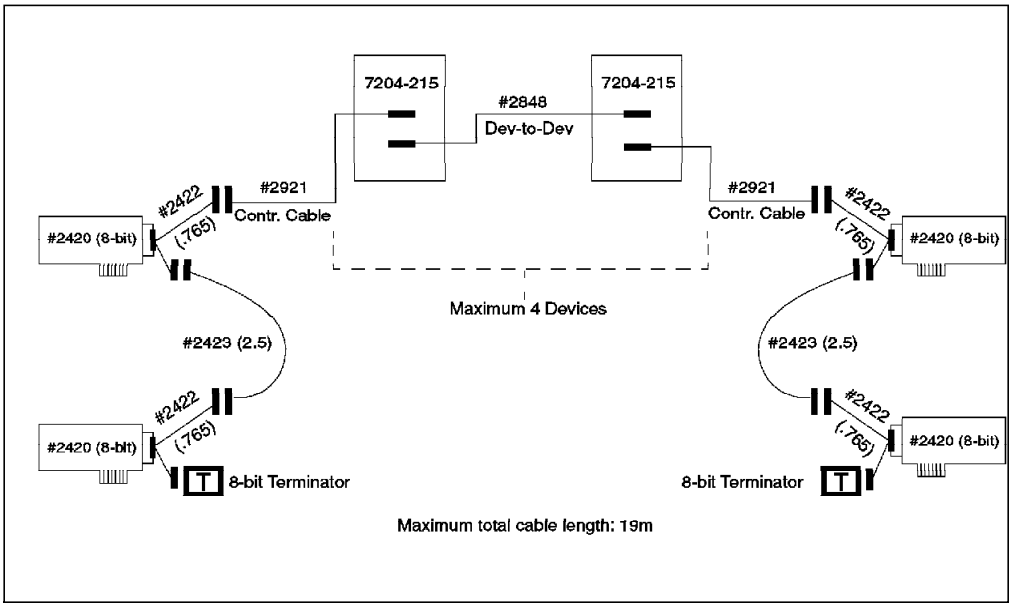

Figure 7. 7204-215 External Disk Drives Connected on an 8-Bit Shared SCSI Bus

D.1.3.2 7204 Model 315, 317, and 325 External Disk Drives

To attach a chain of 7204 Model 315s, 317s, or 325s, or a combination of them to SCSI-2 Differential Fast/Wide Adapter/As or Enhanced SCSI-2 Differential Fast/Wide Adapter/As on a shared 16-bit SCSI bus, you need the following 16-bit cables and terminators:

- 16-Bit SCSI-2 Differential Y-Cable

  FC: 2426 (0.94m), PN: 52G4234

- 16-Bit SCSI-2 Differential System-to-System Cable

  FC: 2424 (0.6m), PN: 52G4291 - OR -

  FC: 2425 (2.5m), PN: 52G4233

  This cable is used only if there are more than two nodes attached to the same shared bus.

- 16-Bit SCSI-2 DE Device-to-Device Cable

  FC: 2845 or 9131 (0.6m), PN: 52G4291 - OR -

  FC: 2846 or 9132 (2.5m), PN: 52G4233

- 16-Bit Terminator

  Included in FC 2426 (Y-Cable), PN: 61G8324

Figure 8 shows four RS/6000s, each represented by one SCSI-2 Differential Fast/Wide Adapter/A, connected on a 16-bit bus to a chain of 7204-315s. The connections would be the same for the 7204-317, and Model 325 drives. You could also substitute the Enhanced SCSI-2 Differential Fast/Wide Adapter/A (feature code 2412) for the SCSI-2 Differential Fast/Wide Adapter/As shown in the figure, if you are running HACMP 4.1 for AIX.

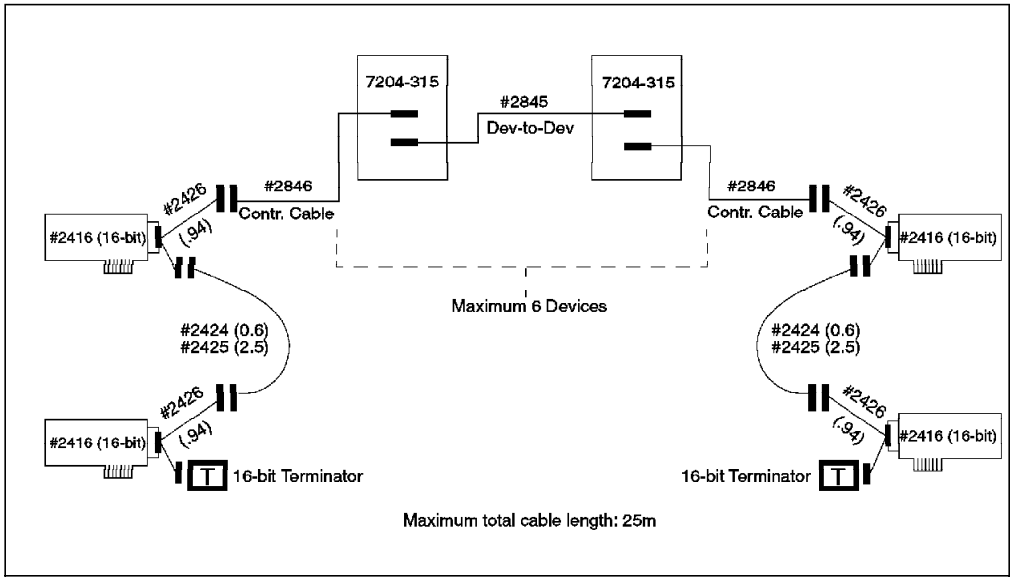

Figure 8. 7204-315 External Disk Drives Connected on a 16-Bit Shared SCSI Bus

D.1.3.3 9334-011 and 9334-501 SCSI Expansion Units

For connecting 9334 Models 011 or 501 to SCSI-2 Differential Controllers on a shared 8-bit SCSI bus, you require the following, in all cases:

- SCSI-2 Differential Y-Cable

  FC: 2422 (0.765m), PN: 52G7348

- SCSI-2 Differential System-to-System Cable

  FC: 2423 (2.5m), PN: 52G7349

  This cable is used only if there are more than two nodes attached to the same shared bus.

- Terminator

  Included in FC 2422 (Y-Cable), PN: 52G7350

In addition to the common set of cables, the 9334-011 requires:

- SCSI-2 DE Controller Cable

  FC: 2921 or 9221 (4.75m), PN: 67G0593 - OR -

  FC: 2923 or 9223 (8.0m), PN: 95X2494

- SCSI-2 DE Device-to-Device Cable

  FC: 2925 or 9225 (2.0m), PN: 95X2492

In addition to the common set of cables, the 9334-501 requires:

- SCSI-2 DE Controller Cable

  FC: 2931 (1.48m), PN: 70F9188 - OR -

  FC: 2933 (2.38m), PN: 45G2858 - OR -

  FC: 2935 (4.75m), PN: 67G0566 - OR -

  FC: 2937 (8.0m), PN: 67G0562

- SCSI-2 DE Device-to-Device Cable:

  FC: 2939 or 9239 (2.0m), PN: 95X2498

Figure 9 on page 114 shows four RS/6000s, each represented by one SCSI-2 Differential Controller, connected on an 8-bit bus to a chain of 9334-011s.
Figure 10 on page 114 shows four RS/6000s, each represented by one SCSI-2 Differential Controller, connected on an 8-bit bus to a chain of 9334-501s.

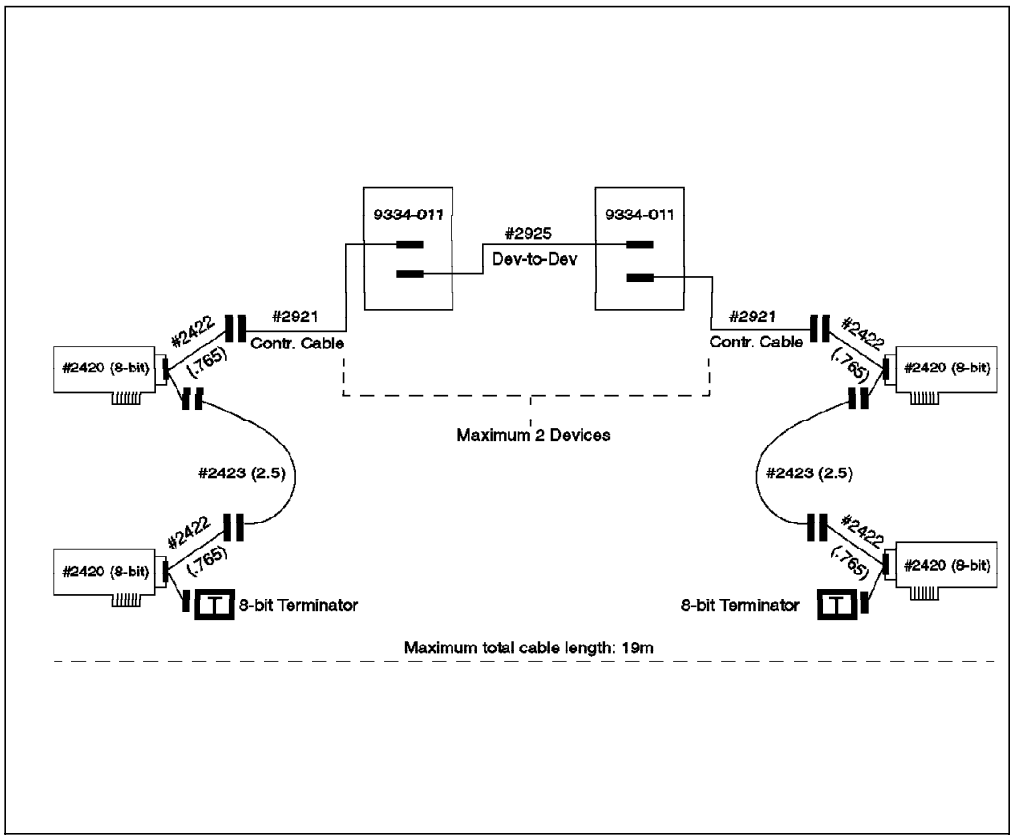

Figure 9. 9334-011 SCSI Expansion Units Connected on an 8-Bit Shared SCSI Bus

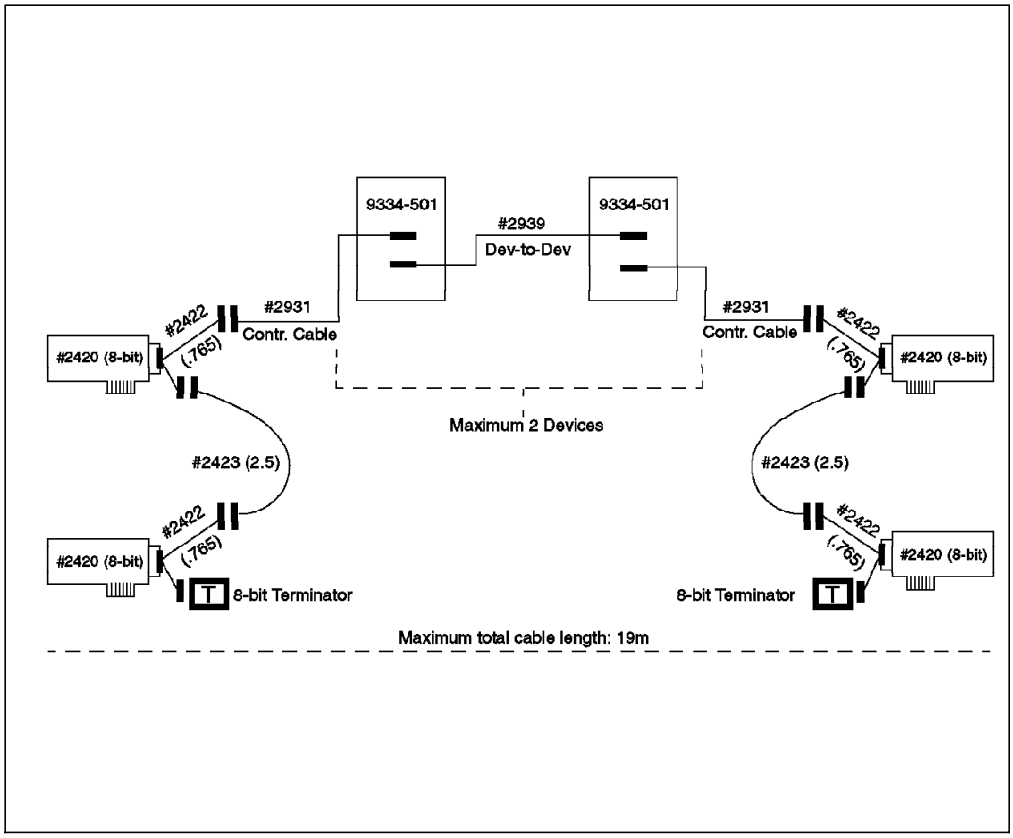

Figure 10. 9334-501 SCSI Expansion Units Connected on an 8-Bit Shared SCSI Bus

D.1.3.4 7134-010 High Density SCSI Disk Subsystem

To attach a 7134-010 to a SCSI-2 Differential Fast/Wide Adapter/A or Enhanced SCSI-2 Differential Fast/Wide Adapter/A on a shared 16-bit SCSI bus, you need the following:

- 16-Bit SCSI-2 Differential Y-Cable

  FC: 2426 (0.94m), PN:52G4234

- 16-Bit SCSI-2 Differential System-to-System Cable

  FC: 2424 (0.6m), PN: 52G4291 - OR -

  FC: 2425 (2.5m), PN: 52G4233

  This cable is used only if there are more than two nodes attached to the same shared bus.

- 16-Bit Differential SCSI Cable

  FC: 2902 (2.4m), PN: 88G5750 - OR -

  FC: 2905 (4.5m), PN: 88G5749 - OR -

  FC: 2912 (12.0m), PN: 88G5747 - OR -

  FC: 2914 (14.0m), PN: 88G5748 - OR -

  FC: 2918 (18.0m), PN: 88G5746

- 16-Bit Terminator (T)

  Included in FC 2426 (Y-Cable), PN: 61G8324

Figure 11 on page 116 shows four RS/6000s, each represented by two SCSI-2 Differential Fast/Wide Adapter/As, connected on a 16-bit bus to a 7134-010 with a base and an expansion unit. You could also substitute the Enhanced SCSI-2 Differential Fast/Wide Adapter/A (feature code 2412) for the SCSI-2 Differential Fast/Wide Adapter/As shown in the figure, if you are running HACMP 4.1 for AIX.

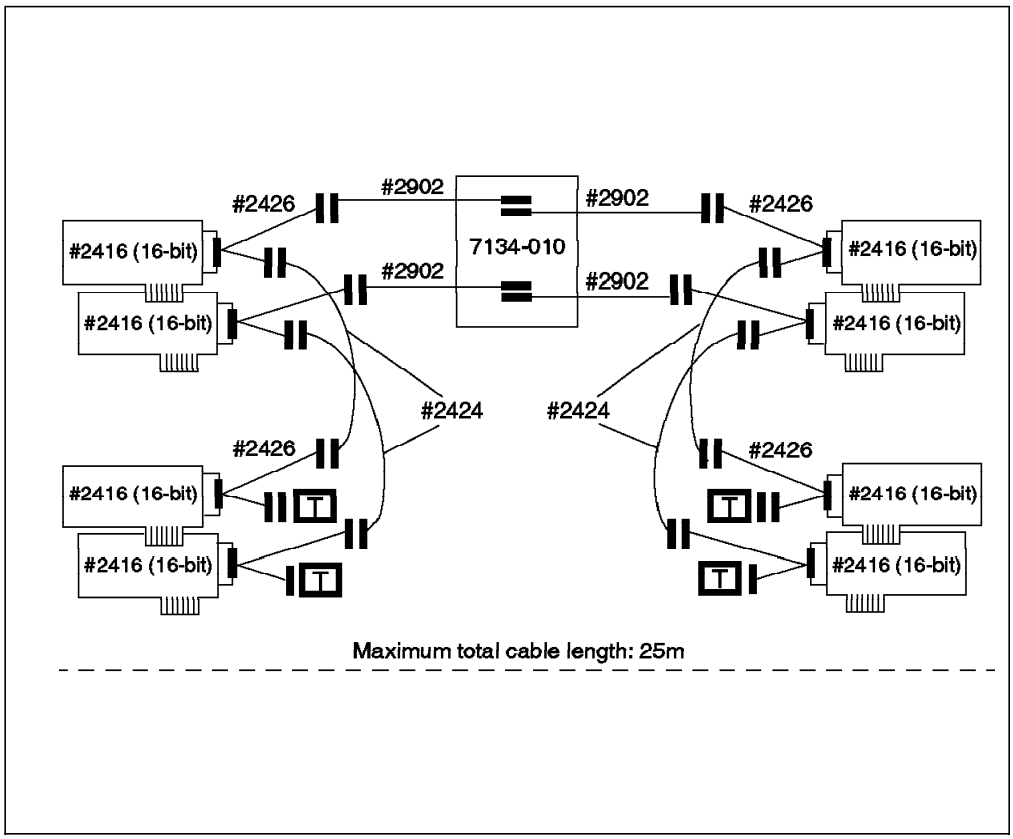

Figure 11. 7134-010 High Density SCSI Disk Subsystem Connected on Two 16-Bit Shared SCSI Buses

D.1.4 AIX's View of Shared SCSI Disks

If your shared SCSI bus has been set up without violating any of the restrictions for termination, SCSI IDs, or cable length, the nodes connected to the shared bus should be able to configure each disk, including the ones inside a 9334 or a 7134, as a separate hdisk device at the next system restart.

D.2 RAID Subsystems

The SCSI adapters that can be used to connect RAID subsystems on a shared SCSI bus in an HACMP cluster are:

- SCSI-2 Differential Controller (FC: 2420, PN: 43G0176)
- SCSI-2 Differential Fast/Wide Adapter/A (FC: 2416, PN: 65G7315)
- Enhanced SCSI-2 Differential Fast/Wide Adapter/A (FC: 2412)

  (This adapter was only supported under AIX 4.1 and HACMP 4.1 for AIX at the time of publishing, but testing was underway to certify the adapter under HACMP/6000 Version 3.1)

The RAID subsystems that you can connect on a shared bus in an HACMP cluster are:

- 7135-110 (HACMP/6000 Version 3.1 only, at the time of publishing) and 7135-210 (HACMP 4.1 for AIX only) RAIDiant Array
- 7137 Model 412, 413, 414, 512, 513, and 514 Disk Array Subsystems

Note: Existing IBM 3514 RAID Array models continue to be supported as shared disk subsystems under HACMP, but since this subsystem has been withdrawn from marketing, it is not described here. As far as cabling and connection characteristics are concerned, the 3514 follows the same rules as the 7137 Disk Array subsystems.

D.2.1 SCSI Adapters

A description of the SCSI adapters that can be used on a shared SCSI bus is given in Section D.1.1, “SCSI Adapters” on page 107.

D.2.2 RAID Enclosures

The 7135 RAIDiant Array can hold a maximum of 30 single-ended disks in two units (one base and one expansion). It has one controller by default, and another controller can be added for improved performance and availability. Each controller takes up one SCSI ID. The disks sit on internal single-ended buses and hence do not take up IDs on the external bus. In an HACMP cluster, each 7135 should have two controllers, each of which is connected to a separate shared SCSI bus. This configuration protects you against any failure (SCSI adapter, cables, or RAID controller) on either SCSI bus.

Because of cable length restrictions, a maximum of two 7135s on a shared SCSI bus is supported by HACMP.

The 7137 Model 412, 413, 414, 512, 513, and 514 Disk Array Subsystems can hold a maximum of eight disks. Each model has one RAID controller, that takes up one SCSI ID on the shared bus. You can have a maximum of two 7137s connected to a maximum of four nodes on an 8-bit or 16-bit shared SCSI bus.

D.2.3 Connecting RAID Subsystems

In this section, we will list the different components required to connect RAID subsystems on a shared bus. We will also show you how to connect these components together.

D.2.3.1 7135-110 or 7135-210 RAIDiant Array

The 7135-110 RAIDiant Array can be connected to multiple systems on either an 8-bit or a 16-bit SCSI-2 differential bus. The Model 210 can only be connected to a 16-bit SCSI-2 Fast/Wide differential bus, using the Enhanced SCSI-2 Differential Fast/Wide Adapter/A.

To connect a set of 7135-110s to SCSI-2 Differential Controllers on a shared 8-bit SCSI bus, you need the following:

- SCSI-2 Differential Y-Cable

  FC: 2422 (0.765m), PN: 52G7348

- SCSI-2 Differential System-to-System Cable

  FC: 2423 (2.5m), PN: 52G7349

  This cable is used only if there are more than two nodes attached to the same shared bus.

- Differential SCSI Cable (RAID Cable)

  FC: 2901 or 9201 (0.6m), PN: 67G1259 - OR -

  FC: 2902 or 9202 (2.4m), PN: 67G1260 - OR -

FC: 2905 or 9205 (4.5m), PN: 67G1261 - OR -

FC: 2912 or 9212 (12m), PN: 67G1262 - OR -

FC: 2914 or 9214 (14m), PN: 67G1263 - OR -

FC: 2918 or 9218 (18m), PN: 67G1264

- Terminator (T)

  Included in FC 2422 (Y-Cable), PN: 52G7350

- Cable Interposer (I)

  FC: 2919, PN: 61G8323

  One of these is required for each connection between a SCSI-2 Differential Y-Cable and a Differential SCSI Cable going to the 7135 unit, as shown in Figure 12.

Figure 12 shows four RS/6000s, each represented by two SCSI-2 Differential Controllers, connected on two 8-bit buses to two 7135-110s each with two controllers.

Note

The diagrams in this book give a logical view of the 7135 subsystem. Please refer to the *7135 Installation and Service Guide* for the exact positions of the controllers and their corresponding connections.

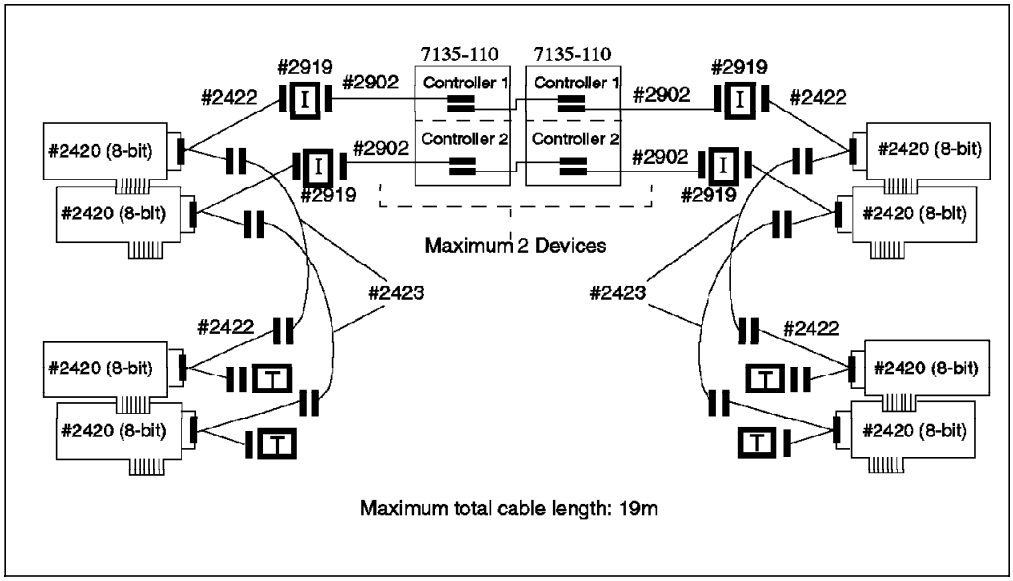

Figure 12. 7135-110 RAIDiant Arrays Connected on Two Shared 8-Bit SCSI Buses

To connect a set of 7135s to SCSI-2 Differential Fast/Wide Adapter/As or Enhanced SCSI-2 Differential Fast/Wide Adapter/As on a shared 16-bit SCSI bus, you need the following:

- 16-Bit SCSI-2 Differential Y-Cable

  FC: 2426 (0.94m), PN: 52G4234

- 16-Bit SCSI-2 Differential System-to-System Cable

  FC: 2424 (0.6m), PN: 52G4291 - OR -

FC: 2425 (2.5m), PN: 52G4233

This cable is used only if there are more than two nodes attached to the same shared bus.

- 16-Bit Differential SCSI Cable (RAID Cable)

  FC: 2901 or 9201 (0.6m), PN: 67G1259 - OR -

  FC: 2902 or 9202 (2.4m), PN: 67G1260 - OR -

  FC: 2905 or 9205 (4.5m), PN: 67G1261 - OR -

  FC: 2912 or 9212 (12m), PN: 67G1262 - OR -

  FC: 2914 or 9214 (14m), PN: 67G1263 - OR -

  FC: 2918 or 9218 (18m), PN: 67G1264

- 16-Bit Terminator (T)

  Included in FC 2426 (Y-Cable), PN: 61G8324

Figure 13 shows four RS/6000s, each represented by two SCSI-2 Differential Fast/Wide Adapter/As, connected on two 16-bit buses to two 7135-110s, each with two controllers.

The 7135-210 requires the Enhanced SCSI-2 Differential Fast/Wide Adapter/A adapter for connection. Other than that, the cabling is exactly the same as shown in Figure 13, if you just substitute the Enhanced SCSI-2 Differential Fast/Wide Adapter/A (FC: 2412) for the SCSI-2 Differential Fast/Wide Adapter/A (FC: 2416) in the picture.

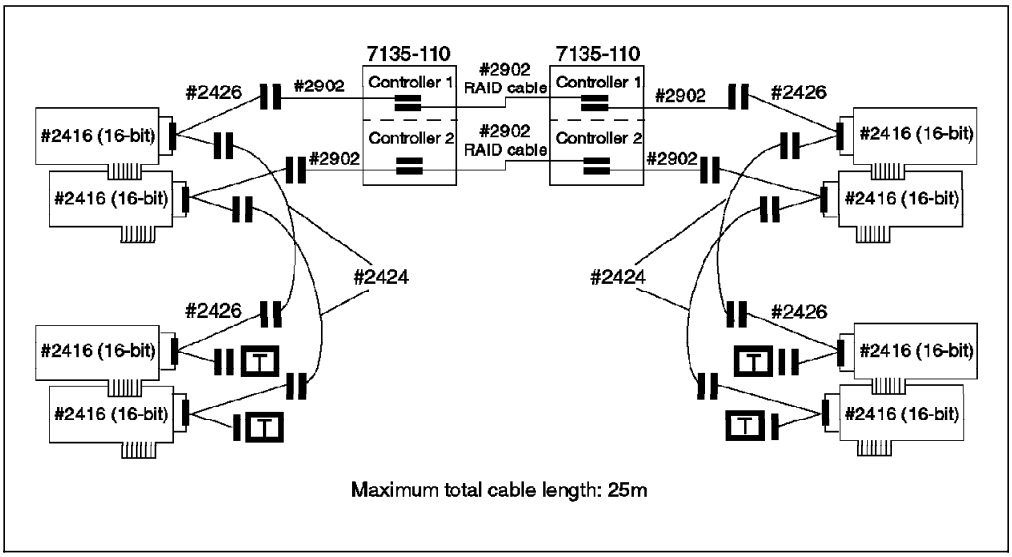

Figure 13. 7135-110 RAIDiant Arrays Connected on Two Shared 16-Bit SCSI Buses

D.2.3.2 7137 Model 412, 413, 414, 512, 513, and 514 Disk Array Subsystems

To connect two 7137s to SCSI-2 Differential Controllers on a shared 8-bit SCSI bus, you need the following:

- SCSI-2 Differential Y-Cable

  FC: 2422 (0.765m), PN: 52G7348

- SCSI-2 Differential System-to-System Cable

  FC: 2423 (2.5m), PN: 52G7349

  This cable is used only if there are more than two nodes attached to the same shared bus.

- Attachment Kit to SCSI-2 Differential High-Performance External I/O Controller

  FC: 2002, PN: 46G4157

  This includes a 4.0-meter cable, an installation diskette, and the *IBM 7137 (or 3514) RISC System/6000 System Attachment Guide*.

- Multiple Attachment Cable

  FC: 3001, PN: 21F9046

  This includes a 2.0-meter cable, an installation diskette, and connection instructions.

- Terminator (T)

  Included in FC 2422 (Y-Cable), PN: 52G7350

Figure 14 shows four RS/6000s, each represented by one SCSI-2 Differential Controller, connected on an 8-bit bus to two 7137s.

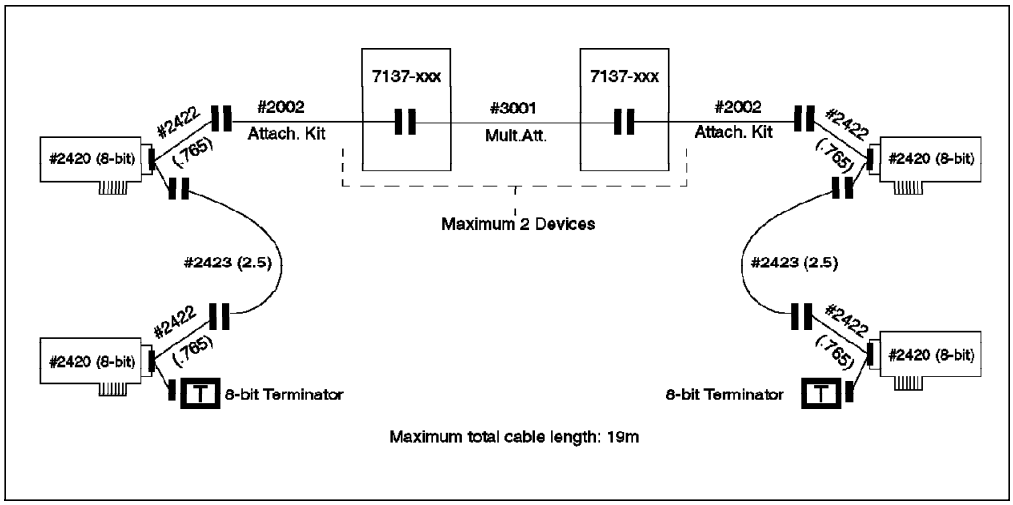

Figure 14. 7137 Disk Array Subsystems Connected on an 8-Bit SCSI Bus

To connect two 7137s to SCSI-2 Differential Fast/Wide Adapter/As or Enhanced SCSI-2 Differential Fast/Wide Adapter/As on a shared 16-bit SCSI bus, you need the following:

- 16-Bit SCSI-2 Differential Y-Cable

  FC: 2426 (0.94m), PN: 52G4234

- 16-Bit SCSI-2 Differential System-to-System Cable

  FC: 2424 (0.6m), PN: 52G4291 - OR -

  FC: 2425 (2.5m), PN: 52G4233

  This cable is used only if there are more than two nodes attached to the same shared bus.

- Attachment Kit to SCSI-2 Differential Fast/Wide Adapter/A or Enhanced SCSI-2 Differential Fast/Wide Adapter/A

  FC: 2014, PN: 75G5028

  This includes a 4.0-meter cable, an installation diskette, and the *IBM 7137 (or 3514) RISC System/6000 System Attachment Guide*.

- Multiple Attachment Cable

  FC: 3001, PN: 21F9046

  This includes a 2.0-meter cable, an installation diskette, and connection instructions.

- 16-Bit Terminator (T)

  Included in FC 2426 (Y-Cable), PN: 61G8324

Figure 15 shows four RS/6000s, each represented by one SCSI-2 Differential Fast/Wide Adapter/As, connected on a 16-bit bus to two 7137s. The Enhanced SCSI-2 Differential Fast/Wide Adapter/A uses exactly the same cabling, and could be substituted for the SCSI-2 Differential Fast/Wide Adapter/A in an AIX 4.1 and HACMP 4.1 for AIX configuration.

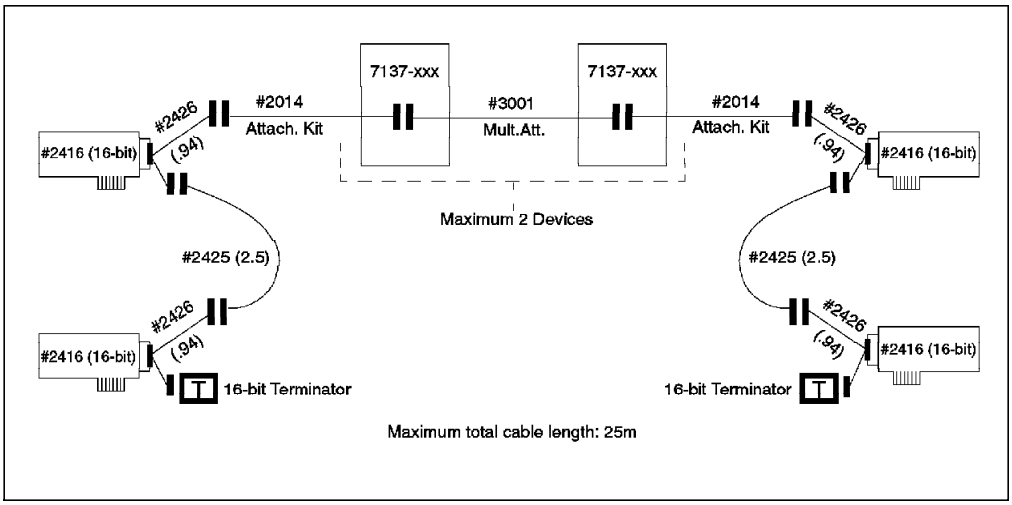

Figure 15. 7137 Disk Array Subsystems Connected on a 16-Bit SCSI Bus

D.2.4 AIX′s View of Shared RAID Devices

The 7135 and 7137 subsystems come preconfigured with Logical Units (LUNs) from the factory. Each LUN gets recognized by nodes on the shared bus as an hdisk device. You can reconfigure the LUNs in a 7135 to suit your requirements by using the 7135 Disk Array Manager software. A 7137 can be reconfigured by using the operator panel on the subsystem itself.

The procedure for configuring LUNs is beyond the scope of this book. Please refer to *7135 RAIDiant Array for AIX - Installation and Reference* for instructions on using the 7135 Disk Array Manager software to create and manage LUNs in a 7135. Please refer to the product documentation that comes with the 7137 subsystem for instructions to set up LUNs on that subsystem.

D.3 Serial Disk Subsystems

To connect serial disk subsystems as shared devices in an HACMP cluster, the adapter that you will use is:

- High-Performance Disk Drive Subsystem Adapter 40/80 MB/sec. (FC: 6212, PN: 67G1755)

The serial disk subsystems that you can connect as shared devices in an HACMP cluster are:

- 9333 Model 011 and 501 High-Performance Disk Drive Subsystems

D.3.1 High-Performance Disk Drive Subsystem Adapter

The High-Performance Disk Drive Subsystem Adapter has four ports, with each port supporting the attachment of a single 9333-011 or 501 controller. Since each controller can drive up to a maximum of four disks, of 2 GB capacity each, you can access up to 32 GB of data with one High-Performance Disk Drive Subsystem Adapter. There is no limit on the number of serial disk adapters that you can have in one node. You do not need to worry about device addresses or terminators with serial disks, since the subsystem is self-addressing. This feature makes it much easier to install and configure than the SCSI options discussed previously.

D.3.2 9333 Disk Subsystems

The 9333 Model 011 and 501 High-Performance Disk Drive Subsystems can each contain a maximum of four disks. The 9333-011 is in a drawer configuration, and is used on rack-mounted models. The 9333-501 is in a mini-tower configuration, and is used on all other models of the RS/6000. Each 9333 subsystem requires a dedicated port on a High-Performance Disk Drive Subsystem Adapter. A maximum of four 9333s can attach to one High-Performance Disk Drive Subsystem Adapter, one for each port. Each 9333 subsystem can be shared with a maximum of eight nodes in a cluster. To connect 9333s to an RS/6000, you need to have AIX Version 3.2.4 or later, and AIX feature 5060 (IBM High-Performance Disk Subsystem Support) installed.

D.3.3 Connecting Serial Disk Subsystems in an HACMP Cluster

To connect a 9333-011 or 501 to two systems, each containing High-Performance Disk Drive Subsystem Adapters, you need the following:

- Serial-Link Cable (Quantity 2)

  FC: 9210 or 3010 (10m)

  FC: 9203 or 3003 (3m)

To connect a 9333-011 or 501 to three or more systems, each containing High-Performance Disk Drive Subsystem Adapters, you need the following:

- Serial-Link Cable (One for each system connection)

  FC: 9210 or 3010 (10m)

  FC: 9203 or 3003 (3m)

- Multiple System Attachment Feature(s)

  FC: 4001 (Connect up to four systems)

FC: 4002 (Connect up to eight systems)

Feature 4001 is a prerequisite for feature 4002.

Figure 16 shows eight RS/6000s, each having a High-Performance Disk Drive Subsystem Adapter, connected to one 9333-501 with the Multiple System Attachment Features 4001 and 4002 installed.

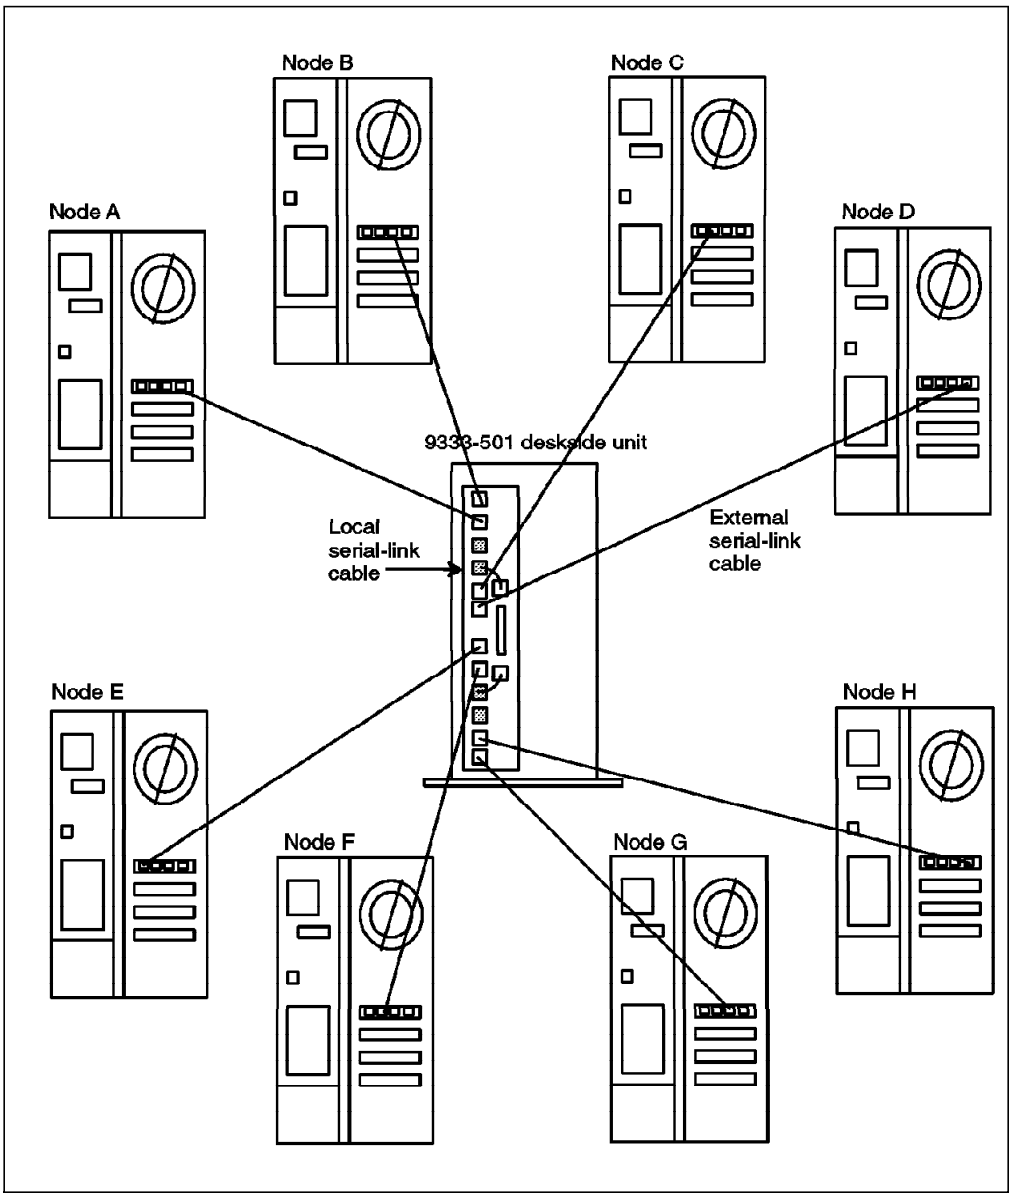

Figure 16. 9333-501 Connected to Eight Nodes in an HACMP Cluster (Rear View)

D.3.4 AIX′s View of Shared Serial Disk Subsystems

Each individual serial disk inside a 9333 subsystem appears as a separate hdisk device on all nodes connected to the subsystem.

D.4 Serial Storage Architecture (SSA) Subsystems

Serial Storage Architecture is a second generation of the high performance serial disk subsystems, started with the IBM 9333 subsystems. SSA subsystems provide new levels of performance, reliability, and flexibility, and are IBM′s strategic high performance disk subsystems for the future.

SSA Support in HACMP

At the time of publishing, the IBM 7133 SSA subsystem was supported for sharing between two nodes only, in a cluster running AIX 3.2.5 and HACMP/6000 Version 3.1. Support for sharing a subsystem between larger numbers of nodes, and support for the the use of the 7133 in an AIX 4.1 and HACMP 4.1 for AIX cluster are expected to be added at a later date. Please check with your IBM representative for the latest support information.

To connect SSA subsystems as shared devices in your HACMP cluster, the adapter that you will use is:

- SSA Four Port Adapter (FC: 6124)

This adapter is shown in Figure 17 on page 125.

The SSA disk subsystems that you can connect as shared devices in an HACMP cluster are:

- IBM 7133-010 SSA Disk Subsystem

  This model is in a drawer configuration, for use in rack mounted systems.

- IBM 7133-500 SSA Disk Subsystem

  This model is in a standalone tower configuration, for use in all models.

D.4.1 SSA Software Requirements

The IBM 7133 SSA Disk Subsystem is supported by AIX Version 3.2.5 with additional program temporary fixes (PTFs), and the AIX 3.2.5 device driver shipped with the SSA Four Port Adapter (FC 6214 on the attaching system). For ease of installation, these PTFs are packaged with the device driver on the CD-ROM shipped with the adapter.

Customers without access to CD-ROM drives on their machines or network can obtain the device driver and required PTFs through the FIXDIST system. The device driver is available as APAR IX52018. The required PTFs, on FIXDIST, are identified as PMP3251.

For alternative delivery, contact your Software Service representative for the appropriate PTFs. The additional Version 3.2.5 PTFs (without the AIX 3.2.5 device driver for the adapter) are included on all AIX Version 3.2.5 orders shipped after May 19, 1995, labelled *AIX 3.2.5 Enhancement 5 (3250-05-00)*.

At the time of publishing, SSA support for AIX 4.1 was expected to be announced by the end of 1995. Please check with your IBM representative for its most current status.

D.4.2 SSA Four Port Adapter

The IBM SSA Four Port Adapter supports the connection of a large capacity of SSA storage. The basic concept of SSA storage connection is that of a loop. An SSA loop starts at one port on the SSA Four Port Adapter continues through a number of SSA disk drives, and concludes at another port on an SSA Four Port Adapter. Each loop can include up to 48 disk devices. Since you can support two loops on each SSA Four Port Adapter, you can support up to 96 disk devices on each adapter. If all those disk devices were of the 4.5 GB capacity, this would provide a potential capacity of 432 GB on an adapter. The adapter itself is shown in Figure 17.

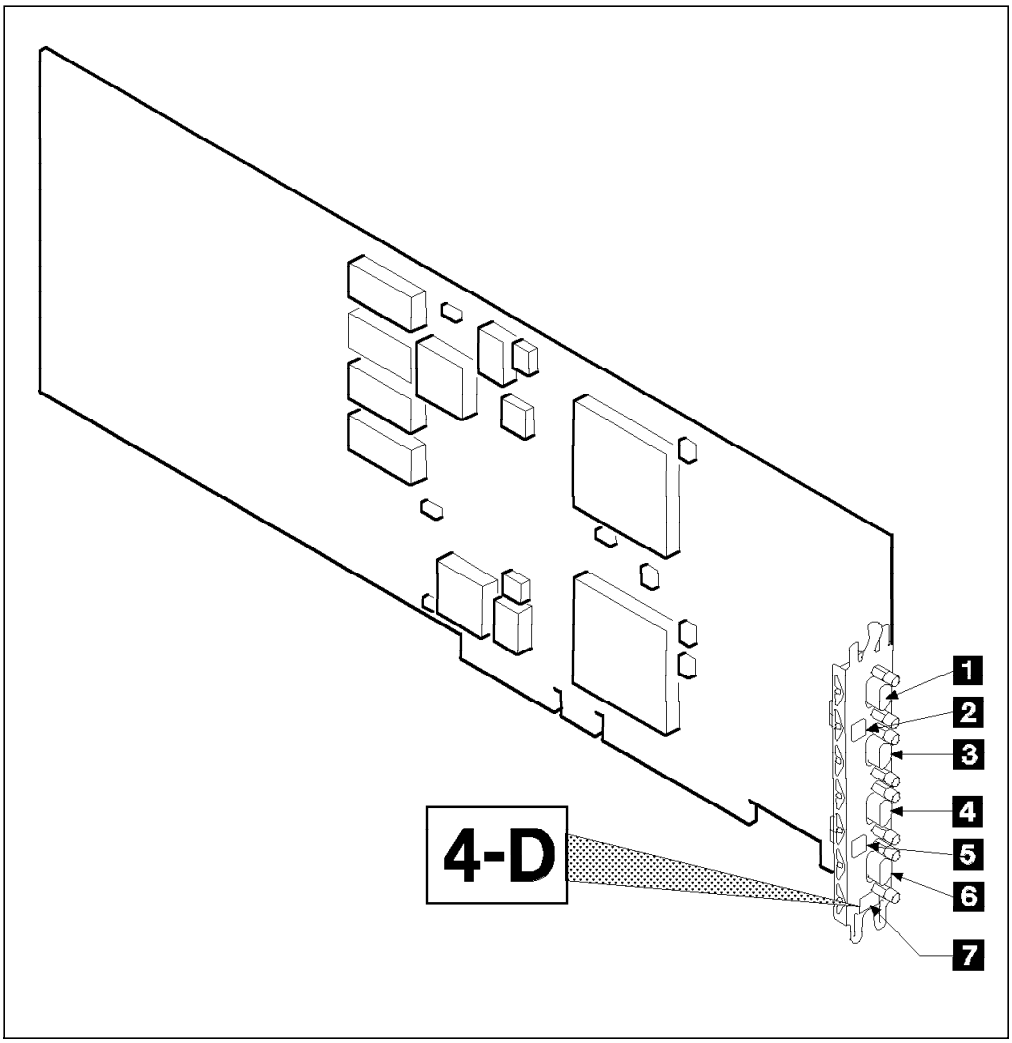

Figure 17. SSA Four Port Aapter

The labeled components of the adapter in the figure are as follows:

1. Connector B2
2. Green light for adapter port pair B
3. Connector B1
4. Connector A2
5. Green light for adapter port pair A
6. Connector A1

7. Type-number label

The green lights for each adapter port pair indicate the status of the attached loop as follows:

Off	Both ports are inactive. If disk drives are connected to these ports, then either the modules have failed or their SSA links have not been enabled.
Permanently on	Both ports are active.
Slow flash	Only one port is active.

The SSA loop that you create need not begin and end on the same &ssaadt.. Loops can be made to go from one adapter to another adapter in the same system or in a different system. There can at most be two adapters on the same loop.

D.4.3 IBM 7133 SSA Disk Subsystem

The IBM 7133 SSA Disk Subsystem is available in two models, the rack drawer model 010 and the standalone tower model 500. While these models hold their disk drives in different physical orientations, they are functionally the same. Each model is capable of holding up to 16 SSA disk drives, each of which can be 1.1 GB, 2.2 GB, or 4.5 GB drives. The subsystem comes standard with four 2.2 GB drives, which can be traded for higher or lower capacity drives at order time.

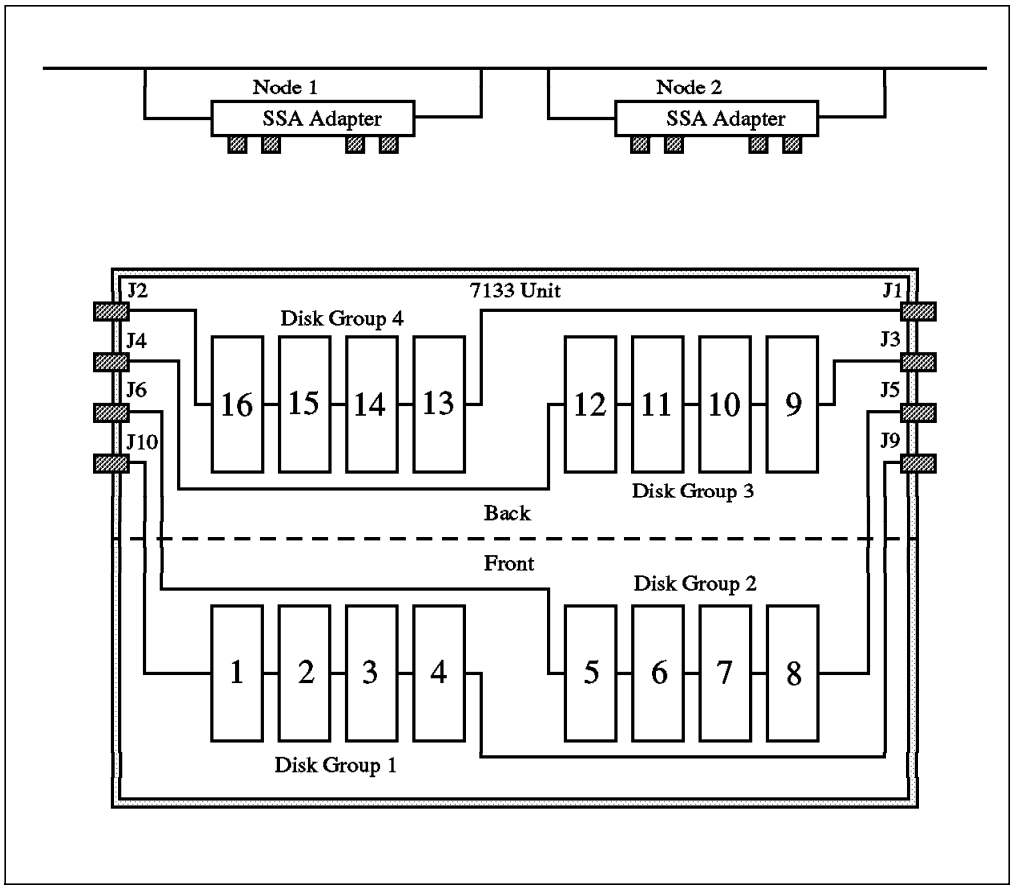

Figure 18. IBM 7133 SSA Disk Subsystem

As you can see in Figure 18, each group of four disk drives in the subsystem is internally cabled as a loop. Disk Group 1 includes disk drive positions 1-4 and is cabled between connectors J9 and J10. Disk Group 2 includes disk drive positions 5-8 and is cabled between connectors J5 and J6. You can also see Disk Groups 3 and 4 in the picture. These internal loops can either be cabled together into larger loops, or individually connected to SSA Four Port Adapters. For instance, if you were to connect a short cable between connectors J6 and J10, you would have a loop of eight drives that could be connected to the SSA Four Port Adapter from connectors J5 and J9.

D.4.4 SSA Cables

SSA cables are available in a variety of different lengths. The connectors at each end are identical, which makes them very easy to use. These cables can be used to connect four disk internal loops together into larger loops within the 7133 subsystem itself, to connect multiple 7133 subsystems together in a larger loop, or to connect a 7133 subsystem to an SSA Four Port Adapter. The same cable can be used for any of these connections, as long as it is long enough. In Table 2 is a list of cable feature codes, along with their lengths, and part numbers:

Table 2. Serial Storage Architecture (SSA) Cables

Cable Description	Feature Code	Part Number
SSA Copper Cable (0.18 meters)	5002	07H9163
SSA Copper Cable (0.6 meters)	5006	31H7960
SSA Copper Cable (1.0 meter)	5010	07H8985
SSA Copper Cable (2.5 meters)	5025	32H1465
SSA Copper Cable (5.0 meters)	5050	88G6406
SSA Copper Cable (10 meters)	5100	32H1466
SSA Copper Cable (25 meters)	5250	88G6406

The feature code numbers start with the number 5, and the next three digits give a rounded length in meters, which makes the feature numbers easy to understand and remember. As was mentioned before, the only difference between these cables is their length. They can be used interchangeably to connect any SSA components together.

If you obtain an announcement letter for the 7133 SSA Subsystem, you will also see a number of other cable feature codes listed, with the same lengths (and same prices) as those in Table 2. You needn't worry or be confused about these, since they are the same cables as those in the tables. As long as you have the correct length of cable for the components you need to connect, you have the right cable.

The maximum distance between components in an SSA loop using IBM cabling is 25 meters. With SSA, there is no special maximum cabling distance for the entire loop. In fact, the maximum cabling distance for the loop would be the maximum distance between components (disks or adapters), mulitplied by the maximum number of components (48) in a loop.

D.4.5 Connecting 7133 SSA Subsystems in an HACMP Cluster

The flexibility of the SSA subsystem creates many different options for attaching SSA subsystems in a cluster, with varying levels of redundancy and availability. Since SSA subsystems are currently only supported for sharing between two nodes, these are the examples that we will use. However, it is expected that you will be able to expand these examples by adding more nodes into the loop(s) in the future. We will illustrate two simple scenarios of SSA connection in this section.

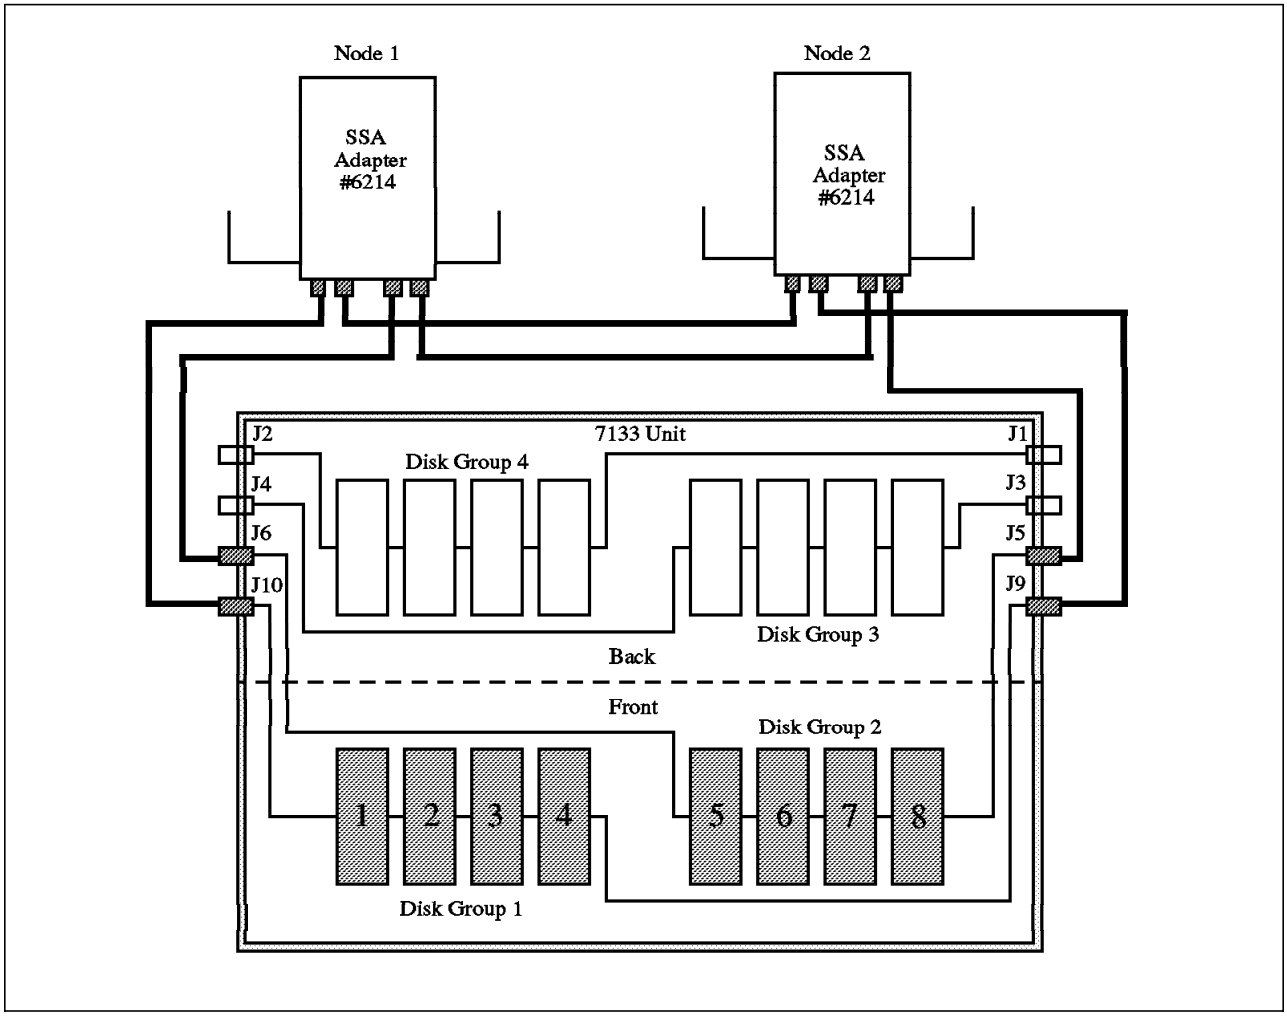

Figure 19. High Availability SSA Cabling Scenario 1

The first scenario, shown in Figure 19, shows a single 7133 subsystem, containing eight disk drives (half full), connected between two nodes in a cluster. We have not labeled the cables, since their lengths will be dependent on the characteristics of your location. Remember, the longest cable currently marketed by IBM is 25 meters, and there are many shorter lengths, as shown in Table 2 on page 127. As we said before, all cables have the same connectors at each end, and therefore are interchangeable, provided they have sufficient length for the task.

In the first scenario, each cluster node has one SSA Four Port Adapter. The disk drives in the 7133 are cabled to the two machines in two loops, the first group of four disks in one loop, and the remaining four in the other. Each of the loops is connected into a different port pair on the SSA Four Port Adapters.

In this configuration, LVM mirroring should be implemented across the two loops; that is, a disk on one loop should be mirrored to a disk on the other loop. Mirroring in this way will protect you against the failure of any single disk drive.

The SSA subsystem is able to deal with any break in the cable in a loop by following the path to a disk in the other direction of the loop, even if it does go through the adapter on the other machine. This recovery is transparent to AIX and HACMP.

The only exposure in this scenario is the failure of one of the SSA Four Port Adapters. In this case, the users on the machine with the failed adapter would lose their access to the disks in the 7133 subsystem. The best solution to this problem is to add a second SSA Four Port Adapter to each node, as shown in Figure 20 on page 130. However, this adds an amount of cost to the solution that might not be justifiable, especially if there is a relatively small amount of disk capacity involved.

An alternative solution would be to use HACMP′s Error Notification feature to protect against the failure. You could define an error notification method, which is triggered on the AIX error log record on the failure of the adapter, and which would run a script to shut down the cluster manager in a *graceful with takeover* mode. This would migrate the users to the other node, from which they would still have access to the disks.

Our second scenario, in Figure 20 on page 130, shows a second SSA Four Port Adapter added to each node. This allows each system to preserve its access to the SSA disks, even if one of the adapters were to fail. This solution does leave an adapter port pair unused on each adapter. These could be used in the future to attach additional loops, if the remaining disk locations in the 7133 were filled, and if additional 7133 subsystems were added into the loops.

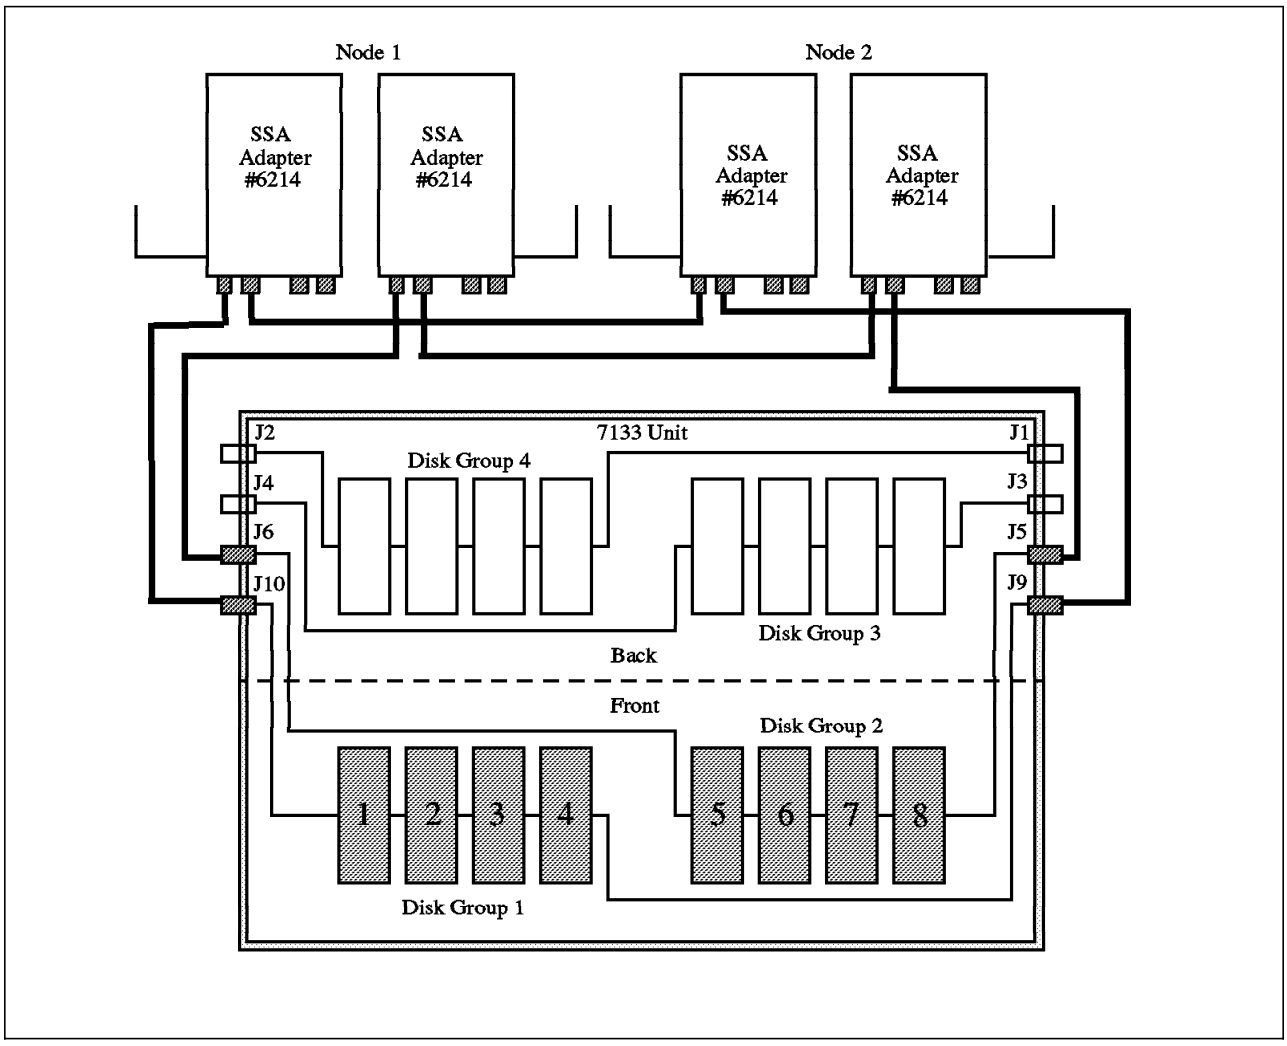

Figure 20. High Availability SSA Cabling Scenario 2

Any of the loops can be extended at any time, by reconnecting the cabling to include the new disks in the loop. If these additions are planned correctly, and cables are unplugged and plugged one at a time, this addition of disks can be done in a “hot-pluggable” way, such that the system does not have to be brought down, access to existing disks is not lost, and the new disks can be configured while the system continues running.

D.4.6 AIX′s View of Shared SSA Disk Subsystems

The AIX operating system configures each disk drive in a shared SSA subsystem as a separate hdisk device on each node.

Appendix E. Example Cluster Planning Worksheets

Cluster Worksheet

Cluster ID 1
Cluster Name disney

Network Name	Network Type	Network Attribute	Netmask	Node Names
trnet1	Token–Ring	public	255.255.255.0	mickey, goofy
etnet1	Ethernet	private	255.255.255.0	mickey, goofy
rsnet1	RS232	serial	N/A	mickey, goofy

Figure 21. Worksheet 1 - Cluster

Network Adapter Worksheet

Node Name mickey

Interface Name	Adapter IP Label	Adapter Function	Adapter IP Address	Network Name	Network Attribute	Adapter HW Address
tr0	mickey	service	9.3.1.79	trnet1	public	42005aa8b484
tr0	mickey_boot	boot	9.3.1.45	trnet1	public	
tr1	mickey_sb	standby	9.3.4.79	trnet1	public	
en0	mickey_en	service	9.3.5.79	etnet1	private	
tty0	mickey_tty0	service	/dev/tty0	rsnet1	serial	

Node Name goofy

Interface Name	Adapter IP Label	Adapter Function	Adapter IP Address	Network Name	Network Attribute	Adapter HW Address
tr0	goofy	service	9.3.1.80	trnet1	public	42005aa8d1f3
tr0	goofy_boot	boot	9.3.1.46	trnet1	public	
tr1	goofy_sb	standby	9.3.4.80	trnet1	public	
en0	goofy_en	service	9.3.5.80	etnet1	private	
tty0	goofy_tty0	service	/dev/tty0	rsnet1	serial	

Figure 22. Worksheet 2 - Network Adapters

9333 Serial Disk Subsystem Configuration Worksheet

Cables Needed: Adapter to Drawer/Tower (Two per drawer/tower unit)

	Node A	Node B	Node C	Node D
Node Name	mickey	goofy	_____	_____
9333 Adapter Label	adapter1	adapter1	_____	_____
Slot Number	3	6	_____	_____
Logical Name	serdasda0	serdasda0	_____	_____
9333 Adapter Label	adapter2	_____	_____	_____
Slot Number	5	_____	_____	_____
Logical Name	serdasda1	_____	_____	_____

9333 Subsyst. Label tower1

	Node A	Node B	Node C	Node D
Adapter I/O Connector	0 (adapter1)	0 (adapter1)	_____	_____
Controller	serdasdc0	serdasdc0	_____	_____

Shared Drives:

Drive	Size	Logical Device Name			
1	857 MB	hdisk1	hdisk2	_____	_____
2	1.07 GB	hdisk2	hdisk3	_____	_____
3	857 MB	hdisk3	hdisk4	_____	_____
4	857 MB	hdisk4	hdisk5	_____	_____

9333 Subsyst. Label tower2

	Node A	Node B	Node C	Node D
Adapter I/O Connector	0 (adapter2)	1 (adapter1)	_____	_____
Controller	serdasdc1	serdasdc1	_____	_____

Shared Drives:

Drive	Size	Logical Device Name			
1	2.0 GB	hdisk5	hdisk6	_____	_____
2	2.0 GB	hdisk6	hdisk7	_____	_____
3	_____	_____	_____	_____	_____
4	_____	_____	_____	_____	_____

Figure 23. Worksheet 3 - 9333 Serial Disk Subsystem Configuration

Shared Volume Group/Filesystem Worksheet

	Node A		Node B	Node C	Node D
Node Names	mickey		goofy	_____	_____
Shared volume group name		test1vg		_____ _____	
Major Number	60		60	_____	_____
Log logical volume name			loglvtest1	_____ _____	
Physical Volumes	hdisk1		hdisk2	_____	_____
	hdisk2		hdisk3	_____	_____

Logical Volume Name	lvtest1
Number of copies per LP	2
On separate PVs?	yes
Filesystem mount point	/test1
Size (MB)	80

Logical Volume Name	_____
Number of copies per LP	_____
On separate PVs?	_____
Filesystem mount point	_____
Size (MB)	_____

Figure 24. Worksheet 4 - Shared Volume Group test1vg

Shared Volume Group/Filesystem Worksheet

	Node A		Node B	Node C	Node D
Node Names	mickey		goofy	_____	_____
Shared volume group name		test2vg		_____	_____
Major Number	61		61	_____	_____
Log logical volume name			loglvtest2	_____	_____
Physical Volumes	hdisk3		hdisk4	_____	_____
	hdisk4		hdisk5	_____	_____

Logical Volume Name	lvtest2
Number of copies per LP	2
On separate PVs?	yes
Filesystem mount point	/test2
Size (MB)	100

Logical Volume Name	_____
Number of copies per LP	_____
On separate PVs?	_____
Filesystem mount point	_____
Size (MB)	_____

Figure 25. Worksheet 5 - Shared Volume Group test2vg

Shared Volume Group/Filesystem Worksheet

	Node A		Node B	Node C	Node D
Node Names	mickey		goofy	_____	_____
Shared volume group name		conc1vg		_____	_____
Major Number	62		62	_____	_____
Log logical volume name		_____		_____	_____
Physical Volumes	hdisk5		hdisk6	_____	_____
	hdisk6		hdisk7	_____	_____

Logical Volume Name	conc1lv
Number of copies per LP	2
On separate PVs?	yes
Filesystem mount point	N/A
Size (MB)	40

Logical Volume Name	conc2lv
Number of copies per LP	2
On separate PVs?	yes
Filesystem mount point	N/A
Size (MB)	28

Figure 26. Worksheet 6 - Shared Volume Group conc1vg

Part 1. Cluster Documentation Tool Report

The following is example output from the documentation tool `doc_dossier` included with this document. A report is produced, in either VM, PostScript, or ascii form, giving detailed configuration information for each node.

You will find that the following is formatted slightly differently from what you will produce on your own system, but it does give you an idea of the information produced.

E.1 Preface of the Report

This document includes:

- All customized files on the system
- System configuration

Its goal is to give a complete picture of a working cluster configuration, including any customizations, at the time it is put into production.

In case of future malfunctions, this will allow the service personnel to understand any changes that have been made to the original cluster configuration.

E.2 SYSTEM CONFIGURATION

E.2.1 Cluster Diagram

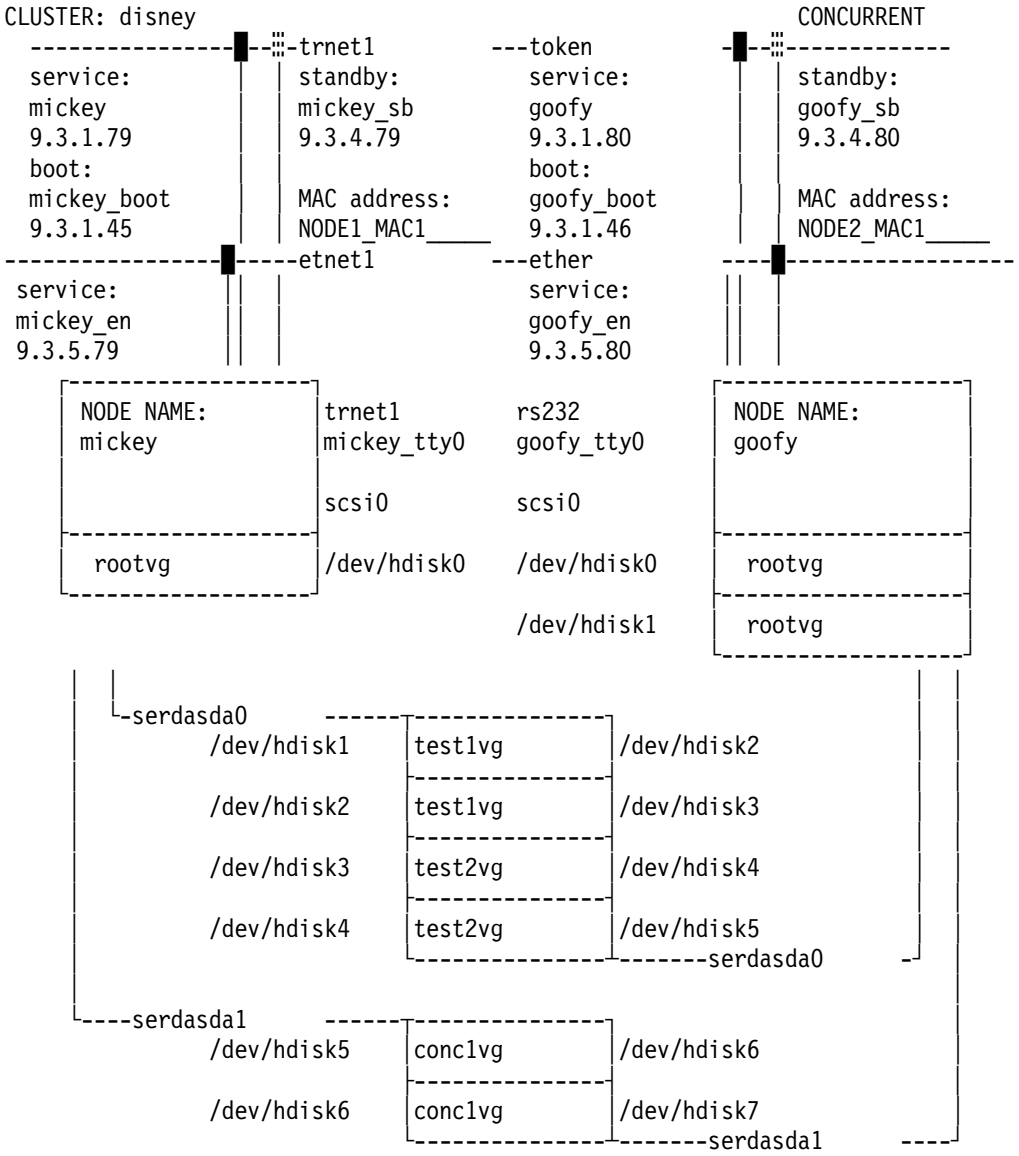

E.2.2 Hostname

```
====> mickey
```

E.2.3 Defined Volume Groups

```
rootvg
test1vg
test2vg
conc1vg
```

```
_____________________________________________
rootvg
VOLUME GROUP:   rootvg                  VG IDENTIFIER:  000147325ccaf23c
VG STATE:       active                  PP SIZE:        4 megabyte(s)
VG PERMISSION:  read/write              TOTAL PPs:      204 (816 megabytes)
MAX LVs:        256                     FREE PPs:       34 (136 megabytes)
LVs:            9                       USED PPs:       170 (680 megabytes)
OPEN LVs:       8                       QUORUM:         2
TOTAL PVs:      1                       VG DESCRIPTORS: 2
STALE PVs:      0                       STALE PPs       0
ACTIVE PVs:     1                       AUTO ON:        yes
_____________________________________________
VOLUME GROUP:   test1vg                 VG IDENTIFIER:  00014732b5a91022
VG STATE:       active                  PP SIZE:        4 megabyte(s)
VG PERMISSION:  read/write              TOTAL PPs:      458 (1832 megabytes)
MAX LVs:        256                     FREE PPs:       416 (1664 megabytes)
LVs:            2                       USED PPs:       42 (168 megabytes)
OPEN LVs:       0                       QUORUM:         1
TOTAL PVs:      2                       VG DESCRIPTORS: 3
STALE PVs:      0                       STALE PPs       0
ACTIVE PVs:     2                       AUTO ON:        no
_____________________________________________
VOLUME GROUP:   test2vg                 VG IDENTIFIER:  00014732ca66234e
VG STATE:       active                  PP SIZE:        4 megabyte(s)
VG PERMISSION:  read/write              TOTAL PPs:      406 (1624 megabytes)
MAX LVs:        256                     FREE PPs:       354 (1416 megabytes)
LVs:            2                       USED PPs:       52 (208 megabytes)
OPEN LVs:       0                       QUORUM:         1
TOTAL PVs:      2                       VG DESCRIPTORS: 3
STALE PVs:      0                       STALE PPs       0
ACTIVE PVs:     2                       AUTO ON:        no
_____________________________________________
VOLUME GROUP:   conc1vg                 VG IDENTIFIER:  00014732b5ac04be
VG STATE:       active                  PP SIZE:        4 megabyte(s)
VG PERMISSION:  read/write              TOTAL PPs:      958 (3832 megabytes)
MAX LVs:        256                     FREE PPs:       924 (3696 megabytes)
LVs:            2                       USED PPs:       34 (136 megabytes)
OPEN LVs:       0                       QUORUM:         2
TOTAL PVs:      2                       VG DESCRIPTORS: 3
STALE PVs:      0                       STALE PPs       0
ACTIVE PVs:     2                       AUTO ON:        no
_____________________________________________
```

E.2.4 Active Volume Groups

```
rootvg
```

E.2.5 Adapters and Disks

```
_____________________________________

scsi0 is a SCSI adapter

The scsi0 adapter has its SCSI ID set to id 7
 and has the following disks connected to it:
ADAPT       DISK            ADDRESS    VOLUME GROUP
scsi0       hdisk0     00-08-00-00     rootvg

The SERIAL adapter serdasda0 has the following disks connected to it:
ADAPT       DISK            ADDRESS    VOLUME GROUP
serdasda0        hdisk1     00-03-00-00     test1vg
serdasda0        hdisk2     00-03-00-01     test1vg
serdasda0        hdisk3     00-03-00-02     test2vg
serdasda0        hdisk4     00-03-00-03     test2vg

The SERIAL adapter serdasda1 has the following disks connected to it:
ADAPT       DISK            ADDRESS    VOLUME GROUP
serdasda1        hdisk5     00-05-00-02     conc1vg
serdasda1        hdisk6     00-05-00-03     conc1vg
_____________________________________

DISK TYPES
hdisk0 857 MB SCSI Disk Drive
hdisk1 857MB Serial-Link Disk Drive
hdisk2 1.07GB Serial-Link Disk Drive
hdisk3 857MB Serial-Link Disk Drive
hdisk4 857MB Serial-Link Disk Drive
hdisk5 2.0GB Serial-Link Disk Drive
hdisk6 2.0GB Serial-Link Disk Drive
```

E.2.6 Physical Volumes

```
_____________________________________
rootvg:
       hdisk0 Available 00-08-00-00 857 MB SCSI Disk Drive
_____________________________________
test1vg:
       hdisk1 Available 00-03-00-00 857MB Serial-Link Disk Drive
       hdisk2 Available 00-03-00-01 1.07GB Serial-Link Disk Drive
_____________________________________
test2vg:
       hdisk3 Available 00-03-00-02 857MB Serial-Link Disk Drive
       hdisk4 Available 00-03-00-03 857MB Serial-Link Disk Drive
_____________________________________
conc1vg:
       hdisk5 Available 00-05-00-02 2.0GB Serial-Link Disk Drive
       hdisk6 Available 00-05-00-03 2.0GB Serial-Link Disk Drive
```

E.2.7 Logical Volumes by Volume Group

```
rootvg
rootvg:
LV NAME             TYPE       LPs   PPs  PVs  LV STATE      MOUNT POINT
hd8                 jfslog     1     1    1    open/syncd    N/A
hd6                 paging     20    20   1    open/syncd    N/A
hd4                 jfs        3     3    1    open/syncd    /
hd1                 jfs        1     1    1    open/syncd    /home
hd3                 jfs        5     5    1    open/syncd    /tmp
hd2                 jfs        135   135  1    open/syncd    /usr
hd9var              jfs        1     1    1    open/syncd    /var
hd5                 boot       2     2    1    closed/syncd  /blv
hd7                 sysdump    2     2    1    open/syncd    /mnt
_____________________________________________________________
test1vg:
LV NAME             TYPE       LPs   PPs  PVs  LV STATE      MOUNT POINT
loglvtest1          jfslog     1     2    2    closed/syncd  N/A
lvtest1             jfs        20    40   2    closed/syncd  /test1
_____________________________________________________________
test2vg:
LV NAME             TYPE       LPs   PPs  PVs  LV STATE      MOUNT POINT
loglvtest2          jfslog     1     2    2    closed/syncd  N/A
lvtest2             jfs        25    50   2    closed/syncd  /test2
_____________________________________________________________
conc1vg:
LV NAME             TYPE       LPs   PPs  PVs  LV STATE      MOUNT POINT
conc1lv             jfs        10    20   2    closed/syncd  N/A
conc2lv             jfs        7     14   2    closed/syncd  N/A
_____________________________________________________________
```

E.2.8 Logical Volume Definitions

```
_____________________________________________________________
LOGICAL VOLUME:     hd8                   VOLUME GROUP:   rootvg
LV IDENTIFIER:      000147325ccaf23c.1    PERMISSION:     read/write
VG STATE:           inactive              LV STATE:       opened/syncd
TYPE:               jfslog                WRITE VERIFY:   off
MAX LPs:            128                   PP SIZE:        4 megabyte(s)
COPIES:             1                     SCHED POLICY:   parallel
LPs:                1                     PPs:            1
STALE PPs:          0                     BB POLICY:      relocatable
INTER-POLICY:       minimum               RELOCATABLE:    yes
INTRA-POLICY:       center                UPPER BOUND     32
MOUNT POINT:        N/A                   LABEL:          None
MIRROR WRITE CONSISTENCY: on
EACH LP COPY ON A SEPARATE PV ?: yes
_____________________________________________________________
LOGICAL VOLUME:     hd6                   VOLUME GROUP:   rootvg
LV IDENTIFIER:      000147325ccaf23c.2    PERMISSION:     read/write
VG STATE:           inactive              LV STATE:       opened/syncd
TYPE:               paging                WRITE VERIFY:   off
MAX LPs:            128                   PP SIZE:        4 megabyte(s)
COPIES:             1                     SCHED POLICY:   parallel
LPs:                20                    PPs:            20
STALE PPs:          0                     BB POLICY:      non-relocatable
```

```
INTER-POLICY:       minimum              RELOCATABLE:    yes
INTRA-POLICY:       middle               UPPER BOUND     32
MOUNT POINT:        N/A                  LABEL:          None
MIRROR WRITE CONSISTENCY: off
EACH LP COPY ON A SEPARATE PV ?: yes

-------------------------------------------------------------------------------
LOGICAL VOLUME:     hd4                  VOLUME GROUP:   rootvg
LV IDENTIFIER:      000147325ccaf23c.3   PERMISSION:     read/write
VG STATE:           inactive             LV STATE:       opened/syncd
TYPE:               jfs                  WRITE VERIFY:   off
MAX LPs:            128                  PP SIZE:        4 megabyte(s)
COPIES:             1                    SCHED POLICY:   parallel
LPs:                3                    PPs:            3
STALE PPs:          0                    BB POLICY:      relocatable
INTER-POLICY:       minimum              RELOCATABLE:    yes
INTRA-POLICY:       center               UPPER BOUND     32
MOUNT POINT:        /                    LABEL:          /
MIRROR WRITE CONSISTENCY: on
EACH LP COPY ON A SEPARATE PV ?: yes

-------------------------------------------------------------------------------
LOGICAL VOLUME:     hd1                  VOLUME GROUP:   rootvg
LV IDENTIFIER:      000147325ccaf23c.4   PERMISSION:     read/write
VG STATE:           inactive             LV STATE:       opened/syncd
TYPE:               jfs                  WRITE VERIFY:   off
MAX LPs:            128                  PP SIZE:        4 megabyte(s)
COPIES:             1                    SCHED POLICY:   parallel
LPs:                1                    PPs:            1
STALE PPs:          0                    BB POLICY:      relocatable
INTER-POLICY:       minimum              RELOCATABLE:    yes
INTRA-POLICY:       center               UPPER BOUND     32
MOUNT POINT:        /home                LABEL:          /home
MIRROR WRITE CONSISTENCY: on
EACH LP COPY ON A SEPARATE PV ?: yes

-------------------------------------------------------------------------------
LOGICAL VOLUME:     hd3                  VOLUME GROUP:   rootvg
LV IDENTIFIER:      000147325ccaf23c.5   PERMISSION:     read/write
VG STATE:           inactive             LV STATE:       opened/syncd
TYPE:               jfs                  WRITE VERIFY:   off
MAX LPs:            128                  PP SIZE:        4 megabyte(s)
COPIES:             1                    SCHED POLICY:   parallel
LPs:                5                    PPs:            5
STALE PPs:          0                    BB POLICY:      relocatable
INTER-POLICY:       minimum              RELOCATABLE:    yes
INTRA-POLICY:       center               UPPER BOUND     32
MOUNT POINT:        /tmp                 LABEL:          /tmp
MIRROR WRITE CONSISTENCY: on
EACH LP COPY ON A SEPARATE PV ?: yes

-------------------------------------------------------------------------------
LOGICAL VOLUME:     hd2                  VOLUME GROUP:   rootvg
LV IDENTIFIER:      000147325ccaf23c.6   PERMISSION:     read/write
VG STATE:           inactive             LV STATE:       opened/syncd
TYPE:               jfs                  WRITE VERIFY:   off
MAX LPs:            512                  PP SIZE:        4 megabyte(s)
COPIES:             1                    SCHED POLICY:   parallel
LPs:                135                  PPs:            135
```

```
STALE PPs:          0                   BB POLICY:      relocatable
INTER-POLICY:       minimum             RELOCATABLE:    yes
INTRA-POLICY:       center              UPPER BOUND     32
MOUNT POINT:        /usr                LABEL:          /usr
MIRROR WRITE CONSISTENCY: on
EACH LP COPY ON A SEPARATE PV ?: yes

-------------------------------------------------------------------------------
LOGICAL VOLUME:     hd9var              VOLUME GROUP:   rootvg
LV IDENTIFIER:      000147325ccaf23c.7  PERMISSION:     read/write
VG STATE:           inactive            LV STATE:       opened/syncd
TYPE:               jfs                 WRITE VERIFY:   off
MAX LPs:            128                 PP SIZE:        4 megabyte(s)
COPIES:             1                   SCHED POLICY:   parallel
LPs:                1                   PPs:            1
STALE PPs:          0                   BB POLICY:      relocatable
INTER-POLICY:       minimum             RELOCATABLE:    yes
INTRA-POLICY:       center              UPPER BOUND     32
MOUNT POINT:        /var                LABEL:          /var
MIRROR WRITE CONSISTENCY: on
EACH LP COPY ON A SEPARATE PV ?: yes

-------------------------------------------------------------------------------
LOGICAL VOLUME:     hd5                 VOLUME GROUP:   rootvg
LV IDENTIFIER:      000147325ccaf23c.8  PERMISSION:     read/write
VG STATE:           inactive            LV STATE:       closed/syncd
TYPE:               boot                WRITE VERIFY:   off
MAX LPs:            128                 PP SIZE:        4 megabyte(s)
COPIES:             1                   SCHED POLICY:   parallel
LPs:                2                   PPs:            2
STALE PPs:          0                   BB POLICY:      relocatable
INTER-POLICY:       minimum             RELOCATABLE:    no
INTRA-POLICY:       edge                UPPER BOUND     32
MOUNT POINT:        /blv                LABEL:          None
MIRROR WRITE CONSISTENCY: on
EACH LP COPY ON A SEPARATE PV ?: yes

-------------------------------------------------------------------------------
LOGICAL VOLUME:     hd7                 VOLUME GROUP:   rootvg
LV IDENTIFIER:      000147325ccaf23c.9  PERMISSION:     read/write
VG STATE:           inactive            LV STATE:       opened/syncd
TYPE:               sysdump             WRITE VERIFY:   off
MAX LPs:            128                 PP SIZE:        4 megabyte(s)
COPIES:             1                   SCHED POLICY:   parallel
LPs:                2                   PPs:            2
STALE PPs:          0                   BB POLICY:      relocatable
INTER-POLICY:       minimum             RELOCATABLE:    yes
INTRA-POLICY:       edge                UPPER BOUND     32
MOUNT POINT:        /mnt                LABEL:          None
MIRROR WRITE CONSISTENCY: on
EACH LP COPY ON A SEPARATE PV ?: yes

-------------------------------------------------------------------------------
LOGICAL VOLUME:     loglvtest1          VOLUME GROUP:   test1vg
LV IDENTIFIER:      00014732b5a91022.1  PERMISSION:     read/write
VG STATE:           inactive            LV STATE:       closed/syncd
TYPE:               jfslog              WRITE VERIFY:   off
MAX LPs:            128                 PP SIZE:        4 megabyte(s)
COPIES:             2                   SCHED POLICY:   parallel
```

```
LPs:                    1                       PPs:            2
STALE PPs:              0                       BB POLICY:      relocatable
INTER-POLICY:           minimum                 RELOCATABLE:    yes
INTRA-POLICY:           center                  UPPER BOUND     32
MOUNT POINT:            N/A                     LABEL:          None
MIRROR WRITE CONSISTENCY: on
EACH LP COPY ON A SEPARATE PV ?: yes

-------------------------------------------------------------------------------
LOGICAL VOLUME:         lvtest1                 VOLUME GROUP:   test1vg
LV IDENTIFIER:          00014732b5a91022.2      PERMISSION:     read/write
VG STATE:               inactive                LV STATE:       closed/syncd
TYPE:                   jfs                     WRITE VERIFY:   off
MAX LPs:                128                     PP SIZE:        4 megabyte(s)
COPIES:                 2                       SCHED POLICY:   parallel
LPs:                    20                      PPs:            40
STALE PPs:              0                       BB POLICY:      relocatable
INTER-POLICY:           minimum                 RELOCATABLE:    yes
INTRA-POLICY:           middle                  UPPER BOUND     32
MOUNT POINT:            /test1                  LABEL:          /test1
MIRROR WRITE CONSISTENCY: on
EACH LP COPY ON A SEPARATE PV ?: yes

-------------------------------------------------------------------------------
LOGICAL VOLUME:         loglvtest2              VOLUME GROUP:   test2vg
LV IDENTIFIER:          00014732ca66234e.1      PERMISSION:     read/write
VG STATE:               inactive                LV STATE:       closed/syncd
TYPE:                   jfslog                  WRITE VERIFY:   off
MAX LPs:                128                     PP SIZE:        4 megabyte(s)
COPIES:                 2                       SCHED POLICY:   parallel
LPs:                    1                       PPs:            2
STALE PPs:              0                       BB POLICY:      relocatable
INTER-POLICY:           minimum                 RELOCATABLE:    yes
INTRA-POLICY:           middle                  UPPER BOUND     32
MOUNT POINT:            N/A                     LABEL:          None
MIRROR WRITE CONSISTENCY: on
EACH LP COPY ON A SEPARATE PV ?: yes

-------------------------------------------------------------------------------
LOGICAL VOLUME:         lvtest2                 VOLUME GROUP:   test2vg
LV IDENTIFIER:          00014732ca66234e.2      PERMISSION:     read/write
VG STATE:               inactive                LV STATE:       closed/syncd
TYPE:                   jfs                     WRITE VERIFY:   off
MAX LPs:                128                     PP SIZE:        4 megabyte(s)
COPIES:                 2                       SCHED POLICY:   parallel
LPs:                    25                      PPs:            50
STALE PPs:              0                       BB POLICY:      relocatable
INTER-POLICY:           minimum                 RELOCATABLE:    yes
INTRA-POLICY:           middle                  UPPER BOUND     32
MOUNT POINT:            /test2                  LABEL:          /test2
MIRROR WRITE CONSISTENCY: on
EACH LP COPY ON A SEPARATE PV ?: yes

-------------------------------------------------------------------------------
LOGICAL VOLUME:         conc1lv                 VOLUME GROUP:   conc1vg
LV IDENTIFIER:          00014732b5ac04be.1      PERMISSION:     read/write
VG STATE:               inactive                LV STATE:       closed/syncd
TYPE:                   jfs                     WRITE VERIFY:   off
MAX LPs:                128                     PP SIZE:        4 megabyte(s)
```

```
COPIES:             2                   SCHED POLICY:   parallel
LPs:                10                  PPs:            20
STALE PPs:          0                   BB POLICY:      relocatable
INTER-POLICY:       minimum             RELOCATABLE:    yes
INTRA-POLICY:       center              UPPER BOUND     32
MOUNT POINT:        N/A                 LABEL:          None
MIRROR WRITE CONSISTENCY: off
EACH LP COPY ON A SEPARATE PV ?: yes

_____________________________________________
LOGICAL VOLUME:     conc2lv             VOLUME GROUP:   conc1vg
LV IDENTIFIER:      00014732b5ac04be.2  PERMISSION:     read/write
VG STATE:           inactive            LV STATE:       closed/syncd
TYPE:               jfs                 WRITE VERIFY:   off
MAX LPs:            128                 PP SIZE:        4 megabyte(s)
COPIES:             2                   SCHED POLICY:   parallel
LPs:                7                   PPs:            14
STALE PPs:          0                   BB POLICY:      relocatable
INTER-POLICY:       minimum             RELOCATABLE:    yes
INTRA-POLICY:       center              UPPER BOUND     32
MOUNT POINT:        N/A                 LABEL:          None
MIRROR WRITE CONSISTENCY: off
EACH LP COPY ON A SEPARATE PV ?: yes
```

E.2.9 Filesystems

Name	Nodename	Mount Pt	VFS	Size	Options	Auto	Accounting
/dev/hd4	--	/	jfs	24576	--	yes	no
/dev/hd1	--	/home	jfs	8192	--	yes	no
/dev/hd2	--	/usr	jfs	1105920	--	yes	no
/dev/hd9var	--	/var	jfs	8192	--	yes	no
/dev/hd3	--	/tmp	jfs	40960	--	yes	no
/dev/hd7	--	/mnt	jfs	--	--	no	no
/dev/hd5	--	/blv	jfs	--	--	no	no
/usr/bin/blv.fs	--	/usr/bin/blv.fs	--	--	--	no	no
/dev/extlv1	--	/inst	jfs	--	rw	no	no
/dev/lvtest1	--	/test1	jfs	--	rw	no	no
/dev/lvtest2	--	/test2	jfs	--	rw	no	no

E.2.10 Paging Spaces

Page Space	Physical Volume	Volume Group	Size	%Used	Active	Auto	Type
hd6	hdisk0	rootvg	80MB	25	yes	yes	lv

E.2.11 TCP/IP Parameters

```
lo0: flags=b<UP,BROADCAST,LOOPBACK>
     inet 127.0.0.1 netmask 0xff000000 broadcast 127.255.255.255
en0: flags=2000063<UP,BROADCAST,NOTRAILERS,RUNNING,NOECHO>
     inet 9.3.5.79 netmask 0xffffff00 broadcast 9.3.5.255
en1: flags=2000062<BROADCAST,NOTRAILERS,RUNNING,NOECHO>
et0: flags=2000002<BROADCAST,NOECHO>
et1: flags=2000002<BROADCAST,NOECHO>
tr0: flags=8063<UP,BROADCAST,NOTRAILERS,RUNNING,ALLCAST>
     inet 9.3.1.45 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
_____________________________
Routing tables
Destination      Gateway            Flags   Refcnt Use       Interface
Netmasks:
(root node)
(0)0 ff00 0
(0)0 ffff ff00 0
(root node)

Route Tree for Protocol Family 2:
(root node)
default          itsorusi.itsc.aust UG          2    15781  tr0
9.3.1            mickey_boot.itsc.a U           3    22533  tr0
9.3.4            mickey_sb.itsc.aus U           1   578797  tr1
9.3.5            mickey_en.itsc.aus U           4   671431  en0
127              localhost          U           0   278190  lo0
(root node)

Route Tree for Protocol Family 6:
(root node)
(root node)
Name  Mtu   Network     Address         Ipkts    Ierrs Opkts    Oerrs Coll
lo0   1536  <Link>                      279124     0   279124     0     0
lo0   1536  127         localhost       279124     0   279124     0     0
en0   1500  <Link>                      672530     0   672438     0     0
en0   1500  9.3.5       mickey_en.itsc.  672530     0   672438     0     0
en1*  1500  <Link>                         235     0        0     0     0
et0*  1492  <Link>                           0     0        0     0     0
et1*  1492  <Link>                           0     0        0     0     0
tr1   1492  <Link>                      748576     0   578803     0     0
tr1   1492  9.3.4       mickey_sb.itsc.  748576     0   578803     0     0
tr0   1492  <Link>                       71366     0    38425     0     0
tr0   1492  9.3.1       mickey_boot.its   71366     0    38425     0     0
nameserver      9.3.1.74
domain     itsc.austin.ibm.com
```

E.2.12 NFS: Exported Filesystems

E.2.13 NFS: Mounted Filesystems

```
Name            Nodename   Mount Pt              VFS   Size    Options    Auto Accounting
```

E.2.14 NFS: Other Parameters

```
Slave servers for the domain

Domains that are being served

These NIS daemons will be started.
_____________________________________
```

E.2.15 Daemons and Processes

```
/etc/cron
/etc/inetd
/etc/init
/etc/methods/sdd serdasda0 00000002
/etc/methods/sdd serdasda1 00000002
/etc/qdaemon
/etc/srcmstr
/etc/syncd 60
/etc/syslogd
/etc/uprintfd
/etc/writesrv
/usr/etc/biod 6
/usr/etc/nfsd 8
/usr/etc/portmap
/usr/etc/rpc.lockd
/usr/etc/rpc.mountd
/usr/etc/rpc.statd
/usr/lib/errdemon
/usr/lib/sendmail -bd -q30m
/usr/lpp/info/bin/infod
/usr/sbin/snmpd
clvmd
kproc
swapper
telnetd
```

E.2.16 Subsystems : Status

Subsystem	Group	PID	Status
syslogd	ras	4448	active
portmap	portmap	5987	active
inetd	tcpip	6245	active
infod	infod	7032	active
snmpd	tcpip	9004	active
sendmail	mail	10419	active
biod	nfs	12503	active
rpc.statd	nfs	12254	active
rpc.lockd	nfs	14305	active
qdaemon	spooler	10980	active
writesrv	spooler	7398	active
nfsd	nfs	14721	active
rpc.mountd	nfs	18128	active
lpd	spooler		inoperative
iptrace	tcpip		inoperative
gated	tcpip		inoperative
named	tcpip		inoperative
routed	tcpip		inoperative
rwhod	tcpip		inoperative
timed	tcpip		inoperative
llbd	ncs		inoperative
nrglbd	ncs		inoperative
keyserv	keyserv		inoperative
ypserv	yp		inoperative
ypbind	yp		inoperative

```
ypupdated       yp                      inoperative
yppasswdd       yp                      inoperative
clinfo          cluster                 inoperative
clstrmgr        cluster                 inoperative
cllockd         lock                    inoperative
clsmuxpd        cluster                 inoperative
```

E.2.17 BOS and LPP Installation/Update History

```
Description                                          State     Fix Id
---------------------------------------------------  ------    ------------

X11fnt.coreX.fnt 1.2.3.0
  3250 X11fnt X11-R5 Maintenance Level                  C     U491105
    X11-R5 Core X Fonts                                 C     U435220

X11fnt.ibm850.pc.fnt 1.2.3.0
  3250 X11fnt X11-R5 Maintenance Level                  C     U491105
    X11-R5 IBM PC-850 Fonts                             C     U428079

X11fnt.iso88591.aix.fnt 1.2.3.0
  3250 X11fnt X11-R5 Maintenance Level                  C     U491105

X11fnt.iso88592.fnt 1.2.3.0
  3250 X11fnt X11-R5 Maintenance Level                  C     U491105

X11fnt.iso88593.fnt 1.2.3.0
  3250 X11fnt X11-R5 Maintenance Level                  C     U491105

X11fnt.iso88594.fnt 1.2.3.0
  3250 X11fnt X11-R5 Maintenance Level                  C     U491105

X11fnt.iso88595.fnt 1.2.3.0
  3250 X11fnt X11-R5 Maintenance Level                  C     U491105

X11fnt.iso88597.aix.fnt 1.2.3.0
  3250 X11fnt X11-R5 Maintenance Level                  C     U491105
    AIXwindows Greek (ISO8859-7) Fonts                  C     U411708
    AIXwindows Greek (ISO8859-7) Fonts                  C
    AIXwindows Greek (ISO8859-7) Fonts                  C     U410795
    X11-R5 ISO-8859-7 Fonts                             C     U428080

X11fnt.iso88599.aix.fnt 1.2.3.0
  3250 X11fnt X11-R5 Maintenance Level                  C     U491105
    AIXwindows Turkish (ISO8859-9) Fonts                C     U411708
    AIXwindows Turkish (ISO8859-9) Fonts                C
    AIXwindows Turkish (ISO8859-9) Fonts                C     U410795
    X11-R5 ISO-8859-9 Fonts                             C     U428081

X11fnt.kanji.aix.fnt 1.2.3.0
  3250 X11fnt X11-R5 Maintenance Level                  C     U491105
    AIXwindows Kanji Fonts                              C     U428082
    X11-R5 PC-932 Fonts                                 C     U435221

X11fnt.oldX.fnt 1.2.3.0
```

```
  3250 X11fnt X11-R5 Maintenance Level                    C    U491105

X11rte.ext.obj 1.2.3.0
  3250 X11rte X11-R5 Maintenance Level                    C    U491119
    AIXwindows Run Time Environment Extensions            C    U411705
    AIXwindows Run Time Environment Extensions            C
    AIXwindows Run Time Environment Extensions            C    U409194
    X11-R5 Additional Postscript Fonts                    C    U428192
    X11-R5 Extensions                                     C    U428193
    X11-R5 Info                                           C    U435058
    X11-R5 X Customize Utilities                          C    U435060
    X11-R5 Motif SMIT                                     C    U435062
    X11-R5 X-Desktop                                      C    U435064
    X11-R5 Font Utility                                   C    U435070
    X11-R5 Additional Postscript Utilities                C    U435222

X11rte.motif1.2.obj 1.2.3.0
  3250 X11rte X11-R5 Maintenance Level                    C    U491119
    Motif 1.2 Translated mwmrc Files                      C    U428196
    Motif 1.2 Window Manager Program                      C    U435138

X11rte.obj 1.2.3.0
  3250 X11rte X11-R5 Maintenance Level                    C    U491119
    AIXwindows Run Time Environment                       C    U411705
    AIXwindows Run Time Environment                       C
    AIXwindows Run Time Environment                       C    U409194
    X11-R5 Runtime Environment Fonts                      C    U428198
    X11-R5 Runtime Environment Locales                    C    U428199
    X11-R5 Runtime Environment                            C    U435140
    X11-R5 Runtime Environment Examples                   C    U435223
    X11-R5 Runtime Environment                            C    U436634

bos.data 3.2.0.0
  3250 bos.data Maintenance Level                         C    U491124
    Info Explorer Databases                               C    U435065
    Terminal Capabilities Database                        C    U435118

bos.obj 3.2.0.0
  3250 bos Maintenance Level                              C    U491123
  3251 AIX Maintenance Level                              C    U493251
    Vital User Information                                C    U424153
    Device Diagnostics                                    C    U427865
    POSIX Asynchronous I/O Services                       C    U428206
    User Messaging Utilities                              C    U428212
    ILS Locale Management Utilities                       C    U428215
    C Language Preprocessor                               C    U428218
    Trace Reporting and Error Logging                     C    U428223
    Input Method Library & Keymaps                        C    U428226
    Math Library                                          C    U428231
    Math Library(SYS-V/SAA Error Semantics)               C    U428232
    X10 Library                                           C    U428233
    Trace Reporting Library                               C    U428236
    Network File System                                   C    U428243
    System Resource Controller                            C    U428249
    Base Operating System                                 C    U432415
```

```
Base Operating System                          C    U432416
Base Operating System                          C    U432447
The Base Operating System                      C    U433283
C library - Common Mode                        C    U433342
Network File System                            C    U434427
Trace Reporting and Error Logging              C    U434922
Bourne Shell                                   C    U434992
Korn Shell                                     C    U434993
SYS-V IPC Utilities                            C    U434996
C Library                                      C    U434997
Security Services Library                      C    U434998
Kernel                                         C    U435001
Info Explorer Utilities                        C    U435066
tty Utilities and Device Drivers               C    U435110
Printer Management Utilities                   C    U435111
Spooler Services Library                       C    U435112
Security Related Utilities and Files           C    U435113
System IPL Utilities                           C    U435115
Base Device Drivers                            C    U435116
Device Drivers Reject Utilities                C    U435117
Devices Message Catalog                        C    U435119
Logical Volume Manager                         C    U435120
Diskless Workstation Manager                   C    U435123
System Installation Utilities                  C    U435125
File Archival Utilities                        C    U435126
Device Configuration Utilities                 C    U435127
GXT100/GXT150 Device Drivers                   C    U435155
HFT Utilities and Device Drivers               C    U435156
GXT1000 Device Drivers & Microcode             C    U435157
GT3/4 Family Device Drivers & Microcode        C    U435158
X11-R4 Library                                 C    U435159
Motif 1.1.4 Library                            C    U435160
X11-R4 Toolkit Library                         C    U435161
File Scanning/Searching Utilities              C    U435165
x25 Device Drivers                             C    U435171
Streams Devices, Interfaces & Utilities        C    U435178
Base Network Utilities                         C    U435180
Lan Device Drivers                             C    U435181
Host Communications Device Drivers             C    U435182
Communications Device Drivers                  C    U435184
awk Language Interpreter                       C    U435228
XCOFF File Management Utilities                C    U435229
File Comparison Utilities                      C    U435230
System, Process, Boot Utilities                C    U435231
File Attribute Utilities                       C    U435232
File System Management Utilities               C    U435233
System Accounting                              C    U435234
Data Compression Utilities                     C    U435235
cron Daemon Utilities                          C    U435236
Date & Time Related Utilities                  C    U435237
DIRECTORIES                                    C    U435238
Character Stream Editing Utilities             C    U435239
Maintenance Level Update Utilities             C    U435240
Character Set Tables & Libraries               C    U435241
Device Configuration Library                   C    U435243
```

```
    Curses Standard and Extended Libraries              C    U435244
    Remote Procedure Call Services Library              C    U435245
    Error Logging Utilities                             C    U435246
    Mail Facilities                                     C    U435247
    Man Page Facility                                   C    U435248
    MultiMedia Device Drivers                           C    U435249
    Base NFS Network Utilities                          C    U435250
    Object Data Manager                                 C    U435251
    BSD Disk Quota Utilities                            C    U435252
    Service Information Tool                            C    U435253
    System Management Interface Tool                    C    U435254
    System Activity Reporting                           C    U435255
    Terminal Capability Utilities                       C    U435256
    Video Capture Adapter                               C    U435257
    vi Text Editor                                      C    U435258
    Base Operating System                               C    U435625
    HFT Utilities and Device Drivers                    C    U436256
    POSIX Asynchronous I/O Services                     C    U436267
    tty Utilities and Device Drivers                    C    U436337
    Devices Message Catalog                             C    U436439
    System IPL Utilities                                C    U436739
    Device Drivers Reject Utilities                     C    U436748
    GT3/4 Family Device Drivers & Microcode             C    U436779
    Streams Devices, Interfaces & Utilities             C    U436782
    Base Device Drivers                                 C    U436811
    Application Installation Utilities                  C    U437028
    Object Data Manager                                 C    U437035
    cron Daemon Utilities                               C    U437079
    GXT1000 Device Drivers & Microcode                  C    U437101
    The Base Operating System                           C    U437134
    The Base Operating System                           C    U437135
    The Base Operating System                           C    U437136
    The Base Operating System                           C    U437137
    Communications Device Drivers                       C    U437272
    Device Diagnostics                                  C    U437315
    Kernel                                              C    U437317
    Mail Facilities                                     C    U437398
    3250 Packaging Requisite                            C    U491150

bosadt.bosadt.data 3.2.0.0
  No Maintenance Level Applied.

bosadt.bosadt.obj 3.2.0.0
  3250 bosadt Maintenance Level                         C    U491125
    The bs Program                                      C    U428255
    Locale Management Utilities                         C    U428260
    lex Program                                         C    U428263
    yacc Program                                        C    U428265
    DOS Device Merge Utility                            C    U435121
    Assembler Utilities                                 C    U435259
    C Language Source Utilities                         C    U435260
    FORTRAN Language Utilities                          C    U435261
    lint Program                                        C    U435262
    make Program                                        C    U435263
    Program Debug Utilities                             C    U435264
```

```
    Source Code Control (sccs) Utilities                C    U435265

bosadt.lib.obj 3.2.0.0
  3250 bosadt Maintenance Level                         C    U491125
    BSD System Administration Help                      C    U428266
    HFT Programming Examples                            C    U428267
    New hardware fast library                           C    U428268
    Base Development Libraries & Include files          C    U432448
    Base Development Libraries & Include files          C    U432449
    Base Development Libraries & Include files          C    U432450
    Programming Examples                                C    U435266
    lint Program Rules Databases                        C    U435267
    Include Files                                       C    U435306
    Include Files                                       C    U436252

bosadt.prof.obj 3.2.0.0
  3250 bosadt Maintenance Level                         C    U491125
    Performance Profiling Utilities                     C    U434994
    Base Profiling Support                              C    U435923

bosadt.xde.obj 3.2.0.0
  3250 bosadt Maintenance Level                         C    U491125
    xde Program Debugger                                C    U428269

bosext1.csh.obj 3.2.0.0
  3250 bosext1 Maintenance Level                        C    U491126
    C Shell                                             C    U434995

bosext1.ecs.obj 3.2.0.0
  3250 bosext1 Maintenance Level                        C    U491126

bosext1.extcmds.data 3.2.0.0
  3250 bosext1.data Maintenance Level                   C    U491127
    man Database                                        C    U428271

bosext1.extcmds.obj 3.2.0.0
  3250 bosext1 Maintenance Level                        C    U491126
    Math Calculator Utilities                           C    U428272
    Performance Monitoring Utilities                    C    U435122

bosext1.mh.obj 3.2.0.0
  3250 bosext1 Maintenance Level                        C    U491126
    mh Mail Program                                     C    U435268

bosext1.uucp.obj 3.2.0.0
  3250 bosext1 Maintenance Level                        C    U491126
    uucp Program and Utilities                          C    U435269

bosext1.vdidd.obj 3.2.0.0
  3250 bosext1 Maintenance Level                        C    U491126
    Video Capture Adapter Utilities                     C    U435270

bosext2.acct.obj 3.2.0.0
  3250 bosext2 Maintenance Level                        C    U491128
    System Accounting Utilities                         C    U435271
```

```
bosext2.ate.obj 3.2.0.0
  3250 bosext2 Maintenance Level                C    U491128
    Simple Terminal Emulator                    C    U435272

bosext2.dlc8023.obj 3.2.0.0
  3250 bosext2 Maintenance Level                C    U491128
    8023 Data Link Control                      C    U435172

bosext2.dlcether.obj 3.2.0.0
  3250 bosext2 Maintenance Level                C    U491128
    Ethernet Data Link Control                  C    U435174

bosext2.dlcfddi.obj 3.2.0.0
  3250 bosext2 Maintenance Level                C    U491128
    FDDI Data Link Control                      C    U435173

bosext2.dlcqllc.obj 3.2.0.0
  3250 bosext2 Maintenance Level                C    U491128
    QLLC Data Link Control                      C    U435176

bosext2.dlcsdlc.obj 3.2.0.0
  3250 bosext2 Maintenance Level                C    U491128
    SDLC Data Link Control                      C    U435177

bosext2.dlctoken.obj 3.2.0.0
  3250 bosext2 Maintenance Level                C    U491128
    Token Ring Data Link Control                C    U435175

bosext2.dosutil.obj 3.2.0.0
  3250 bosext2 Maintenance Level                C    U491128
    DOS File & Disk Utilities                   C    U435124

bosext2.games.obj 3.2.0.0
  3250 bosext2 Maintenance Level                C    U491128
    Miscellaneous Amusements                    C    U428284

bosext2.lrn.data 3.2.0.0
  3250 bosext2.data Maintenance Level           C    U491129

bosext2.x25app.obj 3.2.0.0
  3250 bosext2 Maintenance Level                C    U491128
    X25 Applications                            C    U435179

bosnet.ncs.obj 3.2.0.0
  3250 bosnet Maintenance Level                 C    U491130
    Network Computing Services                  C    U428286
    Network Computing Services                  C    U434753

bosnet.nfs.obj 3.2.0.0
  3250 bosnet Maintenance Level                 C    U491130
    NFS Client Utilities                        C    U428287
    NFS Server Utilities                        C    U435128
    NFS SMIT Utilities                          C    U435273
    NFS Client Utilities                        C    U436325
```

```
bosnet.snmpd.obj 3.2.0.0
  3250 bosnet Maintenance Level                          C    U491130
    Simple Network Management Protocol Daemon (Agent)    C    U432417
    SNMP Daemon                                          C    U435274

bosnet.tcpip.obj 3.2.0.0
  3250 bosnet Maintenance Level                          C    U491130
    TCP/IP Client Utilities                              C    U435114
    TCP/IP Server Utilities                              C    U435275
    TCP/IP SMIT Utilities                                C    U435276

bsl.en_US.aix.loc 3.2.0.0
  3250 bsl Maintenance Level                             A    U491131

bsl.en_US.pc.loc 3.2.0.0
  3250 bsl Maintenance Level                             A    U491131

bsmEn_US.msg 3.2.0.0
  3250 bsmEn_US Maintenance Level                        C    U491133
    Base System Messages - U.S. English                  C    U428303
    SMIT Install Messages - U.S. English                 C    U428304
    Base System Messages - U.S. English                  C    U437316

bspiEn_US.info 3.2.5.0
  No Maintenance Level Applied.

bssiEn_US.info 3.2.5.0
  No Maintenance Level Applied.

cluster.client 3.1.0.0
  No Maintenance Level Applied.
    HACMP/6000                                           C    U438726

cluster.clvm 3.1.0.0
  No Maintenance Level Applied.

cluster.server 3.1.0.0
  No Maintenance Level Applied.
    HACMP/6000                                           C    U438726

sd6k_clnt.obj 2.3.0.11
  No Maintenance Level Applied.

serdasd.mc 3.2.0.16
  No Maintenance Level Applied.

sysback.obj 3.2.0.30
  No Maintenance Level Applied.

txtfmt.bib.data 3.2.0.0
  3250 txtfmt.data Maintenance Level                     C    U491156

txtfmt.bib.obj 3.2.0.0
  3250 txtfmt Maintenance Level                          C    U491155
```

```
    Text Formating Bibliography Utilities                C    U428350

 txtfmt.graf.obj 3.2.0.0
   3250 txtfmt Maintenance Level                          C    U491155
     Tektronics Terminal Drivers                          C    U428351

 txtfmt.hplj.fnt 3.2.0.0
   3250 txtfmt Maintenance Level                          C    U491155

 txtfmt.ibm3812.fnt 3.2.0.0
   3250 txtfmt Maintenance Level                          C    U491155
     IBM-3812 Fonts                                       C    U428390

 txtfmt.ibm3816.fnt 3.2.0.0
   3250 txtfmt Maintenance Level                          C    U491155

 txtfmt.spell.data 3.2.0.0
   3250 txtfmt.data Maintenance Level                     C    U491156

 txtfmt.spell.obj 3.2.0.0
   3250 txtfmt Maintenance Level                          C    U491155
     Spell Checker Utilities                              C    U428352

 txtfmt.tfs.data 3.2.0.0
   3250 txtfmt.data Maintenance Level                     C    U491156

 txtfmt.tfs.obj 3.2.0.0
   3250 txtfmt Maintenance Level                          C    U491155
     Text Formatting Utilities                            C    U428353
     Text Formatting Utilities                            C    U439408

 txtfmt.ts.obj 3.2.0.0
   3250 txtfmt Maintenance Level                          C    U491155
     Postscript Formatter                                 C    U428354

 txtfmt.xpv.obj 3.2.0.0
   3250 txtfmt Maintenance Level                          C    U491155
     X Preview Utility                                    C    U435300

 xlccmp.obj 1.3.0.0
   3250 xlccmp 1.3 Maintenance Level                      C    U491204


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.
```

E.2.18 TTY: Definitions

```
_____________________________
tty0:
speed           9600
parity          none
bpc             8
stops           1
xon             yes
term            dumb
login           disable
runmodes        hupcl,cread,brkint,icrnl,opost,tab3,onlcr,isig,icanon,echo,echoe,echok,ech
octl,echoke,imaxbel,iexten
logmodes        hupcl,cread,echoe,cs8,ixon,ixoff
autoconfig      available
imap            none
omap            none
```

E.2.19 ODM: Customized Attributes

```
†bus0†          †bus_iocc_mem†   †0xfffff0†
†conc1lv†       †copies†         †2†
†conc1lv†       †intra†          †c†
†conc1lv†       †lvserial_id†    †00014732b5ac04be.1†
†conc1lv†       †size†           †10†
†conc1vg†       †auto_on†        †n†
†conc1vg†       †pv†             †000009854777a0910000000000000000†
†conc1vg†       †pv†             †000009854777a5c60000000000000000†
†conc1vg†       †state†          †0†
†conc1vg†       †vgserial_id†    †00014732b5ac04be†
†conc2lv†       †copies†         †2†
†conc2lv†       †intra†          †c†
†conc2lv†       †lvserial_id†    †00014732b5ac04be.2†
†conc2lv†       †size†           †7†
†en0†           †broadcast†      †9.3.5.255†
†en0†           †netaddr†        †9.3.5.79†
†en0†           †netmask†        †255.255.255.0†
†en0†           †state†          †up†
†ent0†          †bus_io_addr†    †0x7290†
†ent0†          †dma_bus_mem†    †0x902000†
†ent0†          †dma_lvl†        †0x7†
†ent1†          †bus_intr_lvl†   †0x9†
†ent1†          †bus_mem_addr†   †0xd4000†
†ent1†          †dma_bus_mem†    †0x3a2000†
†ent1†          †dma_lvl†        †0x6†
†gda0†          †bus_mem_start†  †0xc4000†
†gda0†          †dma1_start†     †0xa00000†
†gda0†          †dma2_start†     †0xc00000†
†gda0†          †dma3_start†     †0x1400000†
†gda0†          †dma4_start†     †0x1600000†
†gda0†          †dma_channel†    †0xa†
†gda0†          †int_level†      †0xa†
†hd1†           †intra†          †c†
†hd1†           †label†          †/home†
†hd1†           †lvserial_id†    †000147325ccaf23c.4†
†hd2†           †intra†          †c†
```

```
hd2          label          /usr
hd2          lvserial_id    000147325ccaf23c.6
hd2          size           135
hd3          intra          c
hd3          label          /tmp
hd3          lvserial_id    000147325ccaf23c.5
hd3          size           5
hd4          intra          c
hd4          label          /
hd4          lvserial_id    000147325ccaf23c.3
hd4          size           3
hd5          intra          e
hd5          lvserial_id    000147325ccaf23c.8
hd5          relocatable    n
hd5          size           2
hd5          type           boot
hd6          lvserial_id    000147325ccaf23c.2
hd6          size           20
hd6          type           paging
hd7          intra          e
hd7          lvserial_id    000147325ccaf23c.9
hd7          size           2
hd7          type           sysdump
hd8          intra          c
hd8          lvserial_id    000147325ccaf23c.1
hd8          type           jfslog
hd9var       intra          c
hd9var       label          /var
hd9var       lvserial_id    000147325ccaf23c.7
hdisk0       pvid           000111874109e6740000000000000000
hdisk1       pvid           0000411925a7461000000000000000
hdisk2       pvid           000002992679061e0000000000000000
hdisk3       pvid           00002819699e632f0000000000000000
hdisk4       pvid           000005080b85c6880000000000000000
hdisk5       pvid           000009854777a0910000000000000000
hdisk6       pvid           000009854777a5c600000000000000
hft0         console        1
hft0         default_disp   gda0
hft0         swkb_path      /usr/lib/nls/loc/En_US.hftkeymap
inet0        hostname       mickey
inet0        route          net,,0,9.3.1.74
loglvtest1   copies         2
loglvtest1   intra          c
loglvtest1   lvserial_id    00014732b5a91022.1
loglvtest1   type           jfslog
loglvtest2   copies         2
loglvtest2   lvserial_id    00014732ca66234e.1
loglvtest2   type           jfslog
lvtest1      copies         2
lvtest1      label          /test1
lvtest1      lvserial_id    00014732b5a91022.2
lvtest1      size           20
lvtest2      copies         2
lvtest2      label          /test2
lvtest2      lvserial_id    00014732ca66234e.2
```

```
†lvtest2†       †size†          †25†
†mem0†          †desc†          †32†
†mem0†          †size†          †32†
†mem0†          †type†          †0x8†
†mem1†          †desc†          †32†
†mem1†          †size†          †32†
†mem1†          †type†          †0x8†
†rootvg†        †pv†            †000111874109e67400000000000000000†
†rootvg†        †state†         †0†
†rootvg†        †vgserial_id†   †000147325ccaf23c†
†sa1†           †dma_lvl†       †0x2†
†scsi0†         †ucode†         †/etc/microcode/8d77.32.54†
†serdasda0†     †dma_bus_mem†   †0x250000†
†serdasda0†     †ucode†         †/etc/microcode/8f78.00.16†
†serdasda1†     †bus_intr_lvl†  †0x7†
†serdasda1†     †bus_io_addr†   †0xc400†
†serdasda1†     †dma_bus_mem†   †0x800000†
†serdasda1†     †dma_lvl†       †0x9†
†serdasda1†     †ucode†         †8f78.00.16†
†serdasdc0†     †ucode†         †8f78.00.16†
†serdasdc1†     †desc†          †51†
†serdasdc1†     †ucode†         †8f78.00.16†
†siokb0†        †int_level†     †0x1†
†sys0†          †bootdisk†      †hd5†
†sys0†          †dcache†        †64K†
†sys0†          †icache†        †8K†
†sys0†          †keylock†       †normal†
†sys0†          †modelcode†     †0x0010†
†sys0†          †realmem†       †65536†
†sys0†          †rostime†       †9003071302†
†sys0†          †sysconst†      †/dev/hft†
†test1vg†       †auto_on†       †n†
†test1vg†       †pv†            †000002992679061e0000000000000000†
†test1vg†       †pv†            †0000411925a7461000000000000000000†
†test1vg†       †quorum†        †n†
†test1vg†       †state†         †0†
†test1vg†       †vgserial_id†   †00014732b5a91022†
†test2vg†       †auto_on†       †n†
†test2vg†       †pv†            †000005080b85c6880000000000000000†
†test2vg†       †pv†            †00002819699e632f0000000000000000†
†test2vg†       †quorum†        †n†
†test2vg†       †state†         †0†
†test2vg†       †vgserial_id†   †00014732ca66234e†
†tok0†          †alt_addr†      †0x42005aa8b484†
†tok0†          †dma_bus_mem†   †0x200000†
†tok0†          †ring_speed†    †16†
†tok1†          †alt_addr†      †0x42005aa8d1f3†
†tok1†          †bus_intr_lvl†  †0x5†
†tok1†          †bus_io_addr†   †0x96a0†
†tok1†          †dma_bus_mem†   †0x352000†
†tok1†          †dma_lvl†       †0x5†
†tok1†          †ring_speed†    †16†
†tr0†           †netaddr†       †9.3.1.45†
†tr0†           †netmask†       †255.255.255.0†
†tr0†           †state†         †up†
```

```
†tr1†           †broadcast†     †9.3.4.255†
†tr1†           †netaddr†       †9.3.4.79†
†tr1†           †netmask†       †255.255.255.0†
†tr1†           †state†         †up†
†tty0†          †sttyval†       †3 1c 8 15 4 0 0 11 13 1a 19 12 f 17 16 0 10702 c05 d04bd 2a003b
```

E.3 HACMP CONFIGURATION

E.3.1 Cluster (Command: cllsclstr)

```
ID   Name

1    disney
```

E.3.2 Nodes (Command: cllsnode)

```
NODE goofy:
     Interfaces to network etnet1
          Service Interface: Name goofy_en, Attribute private, IP address 9.3.5.80
     Interfaces to network rsnet1
          Service Interface: Name goofy_tty0, Attribute serial, IP address /dev/tty0
     Interfaces to network trnet1
          Boot Interface: Name goofy_boot, Attribute public, IP address 9.3.1.46
          Service Interface: Name goofy, Attribute public, IP address 9.3.1.80
          Standby Interface: Name goofy_sb, Attribute public, IP address 9.3.4.80

NODE mickey:
     Interfaces to network etnet1
          Service Interface: Name mickey_en, Attribute private, IP address 9.3.5.79
     Interfaces to network rsnet1
          Service Interface: Name mickey_tty0, Attribute serial, IP address /dev/tty0
     Interfaces to network trnet1
          Boot Interface: Name mickey_boot, Attribute public, IP address 9.3.1.45
          Service Interface: Name mickey, Attribute public, IP address 9.3.1.79
          Standby Interface: Name mickey_sb, Attribute public, IP address 9.3.4.79
```

E.3.3 Networks (Command: cllsnw)

```
Network         Attribute  Node       Adapter(s)

etnet1          private    goofy      goofy_en
                           mickey     mickey_en

rsnet1          serial     goofy      goofy_tty0
                           mickey     mickey_tty0

trnet1          public     goofy      (goofy_boot) goofy goofy_sb
                           mickey     (mickey_boot) mickey mickey_sb
```

E.3.4 Adapters (Command: cllsif)

```
Adapter             Type       Network    Net Type   Attribute  Node       IP Address       Hardware Address

goofy_en            service    etnet1     ether      private    goofy      9.3.5.80
goofy_tty0          service    rsnet1     rs232      serial     goofy      /dev/tty0
goofy_boot          boot       trnet1     token      public     goofy      9.3.1.46
goofy               service    trnet1     token      public     goofy      9.3.1.80         0x42005aa8d1f3
goofy_sb            standby    trnet1     token      public     goofy      9.3.4.80
mickey_en           service    etnet1     ether      private    mickey     9.3.5.79
mickey_tty0         service    rsnet1     rs232      serial     mickey     /dev/tty0
```

```
mickey_boot          boot       trnet1     token      public     mickey     9.3.1.45
mickey               service    trnet1     token      public     mickey     9.3.1.79          0x42005aa8b484
mickey_sb            standby    trnet1     token      public     mickey     9.3.4.79
```

E.3.5 Topology (Command: cllscf)

```
Cluster Description of Cluster disney
Cluster ID: 1
There were 3 networks defined : etnet1, rsnet1, trnet1
There are 2 nodes in this cluster.

NODE goofy:
     This node has 3 service interface(s):

     Service Interface goofy_en:
          IP address:     9.3.5.80
          Hardware Address:
          Network:   etnet1
          Attribute:      private

     Service Interface goofy_en has no standby interfaces.


     Service Interface goofy_tty0:
          IP address:     /dev/tty0
          Hardware Address:
          Network:   rsnet1
          Attribute:      serial

     Service Interface goofy_tty0 has no standby interfaces.


     Service Interface goofy:
          IP address:     9.3.1.80
          Hardware Address:    0x42005aa8d1f3
          Network:   trnet1
          Attribute:      public

     Service Interface goofy has a possible boot configuration:
          Boot (Alternate Service) Interface: goofy_boot
          IP address:     9.3.1.46
          Network:   trnet1
          Attribute:      public

     Service Interface goofy has 1 standby interfaces.
          Standby Interface 1: goofy_sb
          IP address:     9.3.4.80
          Network:   trnet1
          Attribute:      public



NODE mickey:
     This node has 3 service interface(s):
```

```
     Service Interface mickey_en:
          IP address:     9.3.5.79
          Hardware Address:
          Network:   etnet1
          Attribute:      private

     Service Interface mickey_en has no standby interfaces.


     Service Interface mickey_tty0:
          IP address:     /dev/tty0
          Hardware Address:
          Network:   rsnet1
          Attribute:      serial

     Service Interface mickey_tty0 has no standby interfaces.


     Service Interface mickey:
          IP address:     9.3.1.79
          Hardware Address:    0x42005aa8b484
          Network:   trnet1
          Attribute:      public

     Service Interface mickey has a possible boot configuration:
          Boot (Alternate Service) Interface: mickey_boot
          IP address:     9.3.1.45
          Network:   trnet1
          Attribute:      public

     Service Interface mickey has 1 standby interfaces.
          Standby Interface 1: mickey_sb
          IP address:     9.3.4.79
          Network:   trnet1
          Attribute:      public




Breakdown of network connections:

Connections to network etnet1
     Node goofy is connected to network etnet1 by these interfaces:
          goofy_en

     Node mickey is connected to network etnet1 by these interfaces:
          mickey_en


Connections to network rsnet1
     Node goofy is connected to network rsnet1 by these interfaces:
          goofy_tty0

     Node mickey is connected to network rsnet1 by these interfaces:
```

```
          mickey_tty0


Connections to network trnet1
     Node goofy is connected to network trnet1 by these interfaces:
          goofy_boot
          goofy
          goofy_sb

     Node mickey is connected to network trnet1 by these interfaces:
          mickey_boot
          mickey
          mickey_sb
```

E.3.6 Resources (Command: clshowres -n All)

```
Run Time Parameters:

Node Name                                   All
Debug Level
Host uses NIS or Name Server
```

E.3.7 Daemons (Command: clshowsrv -a)

```
Subsystem         Group            PID     Status
 clstrmgr         cluster                  inoperative
 clinfo           cluster                  inoperative
 clsmuxpd         cluster                  inoperative
 cllockd          lock                     inoperative
```

E.4 HACMP EVENTS and AIX ERROR NOTIFICATION

- In the following pages you will find shell scripts which have been prefixed by CMD, PRE, POS and REC. Read the explanations given below in order to understand what they are all about.
- When you have understood that, then you will easily understand what they contain.

E.4.1.1 Event Processing Overview

- The HACMP daemons which run on the various cluster nodes all communicate amongst themselves.
- They react to the 32 predefined cluster events such as :
  - Node 2 has just rejoined the cluster
  - A network has just failed
- Default shell scripts for all of the events are in the directory /usr/sbin/cluster/events.
- Some of the scripts are just empty shells which you can customize according to your needs.
- It is advisable NOT to modify the original scripts. Select the event you wish to customize. This is copied into the /usr/HACMP_ANSS/script directory and prefixed by CMD_ (for example, network_down --> CMD_network_down).
- The events are configured in the ODM. The event object class is called /etc/objrepos/HACMPevent. As the location of the event script to be executed is stored within the object, it is necessary to modify the path name, either with SMIT or use the tool and let it do it for you automatically.

E.4.1.2 The PRE and POST shell scripts

- Sometimes it is necessary to carry out a certain action before (PRE) or after (POS) an event script is executed. An example may be sending a message PRE_stop_server before stopping the server application through CMD_stop_server. Then once it has taken place, sending another message via POS_stop_server.
- The PRE and POST events are also modified by SMIT or by the tool.
- They are placed in the /usr/HACMP_ANSS/script directory as well.

E.4.1.3 The RECOVERY shell script

- Each event should send a return code of 0 if it has successfully completed execution. If not, then HACMP will not terminate the event properly and you will see a number of messages on the console.
- We can customize a reaction to a script terminating with a non 0 exit status by executing a RECOVERY script. This script will be executed one or more times depending on how you have set the **Retry Counter** field in the SMIT Event Customization panel.
- Once again the RECOVERY script is configured either through SMIT or with the tool.
- A template is created for you (if you use the tool) in /usr/HACMP_ANSS/script with the event name prefixed by REC_ (for example, REC_network_down). The shell script is empty, and you are free to customize it as you wish.

E.4.1.4 Primary Events

Event	Cause and action
config_too_long	Sends a periodic console message when a node has been in reconfiguration for more than six minutes.

fail_standby	Sends a console message when a standby adapter fails or is no longer available because it has been used to take over the IP address of another adapter.
join_standby	Sends a console message when a standby adapter becomes available.
network_down	Occurs when the cluster determines that a network has failed. The event script provided takes no default action, since the appropriate action will be site/LAN specific.
network_down_complete	Occurs only after a **network_down** event has successfully completed. The event script provided takes no default action, since the appropriate action will be site/LAN specific.
network_up	Occurs when the cluster determines that a network has become available. The event script provided takes no default action, since the appropriate action will be site/LAN specific.
network_up_complete	Occurs only after a **network_up** event has successfully completed. The event script provided takes no default action, since the action will be site/LAN specific.
node_down	Occurs when a node is detaching from the cluster, either voluntarily or due to a failure. Depending on whether the node is local or remote, either the **node_down_local** or **node_down_remote** sub event is called.
node_down_complete	Occurs only after a **node_down** event has successfully completed. Depending on whether the node is local or remote, either the **node_down_local_complete** or **node_down_remote_complete** sub event is called.
node_up	Occurs when a node is joining the cluster. Depending on whether the the node is local or remote, either the **node_up_local** or **node_up_remote** sub event is called.
node_up_complete	Occurs only after a **node_up** event has successfully completed. Depending on whether the node is local or remote, either the **node_up_local_complete** or **node_up_remote_complete** sub event is called.
swap_adapter	Exchanges or swaps the IP addresses of two network interfaces. NIS and name serving are temporarily turned off during this event.
swap_adapter_complete	Occurs only after a **swap_adapter** event has successfully completed. Ensures that the local ARP cache is updated by deleting entries and pinging cluster IP addresses.
event_error	Occurs when an HACMP event script fails for some reason.

E.4.1.5 Secondary Events

Event	Cause and action
acquire_service_addr	Configures boot addresses to the corresponding service address and starts TCP/IP servers and network daemons by running the **telinit -a** command. HACMP modifies the /etc/inittab file by setting all the TCP/IP related startup records to a run level of **a**.
acquire_takeover_addr	Acquires takeover IP address by checking configured standby addresses and swapping them with failed service addresses.
get_disk_vg_fs	Acquire disk, volume group and file system resources as part of takeover.

node_down_local	Releases resources taken from a remote node, stops application servers, releases a service address taken from a remote node, releases concurrent volume groups, unmounts file systems and reconfigures the node to its boot address.
node_down_local_complete	Instructs the cluster manager to exit when the local node has completed detaching from the cluster. This event only occurs after a **node_down_local** event has successfully completed.
node_down_remote	Unmounts any NFS file systems and places a concurrent volume group in non-concurrent mode when the local node is the only surviving node in the cluster. If the failed node did not go down gracefully, acquires a failed node¢s resources: file systems, volume groups and disks and service address.
node_down_remote_complete	Starts takeover application servers if the remote node did not go down gracefully. This event only occurs after a **node_down_remote** event has successfully completed.
node_up_local	When the local node attaches to the cluster: acquires the service address, clears the application server file, acquires file systems, volume groups and disks resources, exports file systems and either activates concurrent volume groups or puts them into concurrent mode depending upon the status of the remote node(s).
node_up_local_complete	Starts application servers and then checks to see if an inactive takeover is needed. This event only occurs after a **node_up_local** event has successfully completed.
node_up_remote	Causes the local node to release all resources taken from the remote node and to place the concurrent volume groups into concurrent mode.
node_up_remote_complete	Allows the local node to do an NFS mount only after the remote node is completely up. This event only occurs after a **node_up_remote** event has successfully completed.
release_service_addr	Detaches the service address and reconfigures to its boot address.
release_takeover_addr	Identifies a takeover address to be released because a standby adapter on the local node is masquerading as the service address of the remote node. Reconfigures the local standby into its original role.
release_vg_fs	Releases volume groups and file systems that the local node took from the remote node.
start_server	Starts application servers.
stop_server	Stops application servers.

E.4.1.6 HARDWARE and SOFTWARE Errors

- AIX has a daemon †errdemon† which is alerted by the kernel whenever a HARDWARE or SOFTWARE incident takes place. Errors are logged into the AIX error log, and can be examined with the †errpt† command.
- There exists an object class /etc/objrepos/errnotify in ODM which can be customized for the special handling of errors. The customization can be carried out with SMIT, and consists of configuring the types of errors to be dealt with, and the action to be taken when such an error occurs. This is done through the definition of a script to be executed when this error is put into the AIX error log.
- The program err_select can also be used for the customization of error handling. It creates templates in /usr/HACMP_ANSS/script for you to customize. All of these templates are prefixed

by error_. The name of the file depends on the type of error selected (for example, error_SCSI).

E.4.2 Script: /usr/HACMP_ANSS/script/CMD_node_down_remote

this file has not been modified

E.4.3 Script: /usr/HACMP_ANSS/script/CMD_node_up_remote

this file has not been modified

E.4.4 Script: /usr/HACMP_ANSS/script/POS_node_down_remote

```
#!/bin/ksh
# program   : POS_node_down_remote
# role      : run after the event
# arguments :   $1 = event name
#                  $2 = return code
# written   : Wed Dec 13 16:43:25 CST 1995
# modified  :
 . /usr/HACMP_ANSS/tools/tool_var
STATUS=0
(print †n=POST-EVENT===============$(date)†
print †on : $(hostname) †
print †AFTER : $1†
print †return code : $2† ) >> $LOG
################################################################################
# Enter your customizing code here

######################## END OF CUSTOMIZATION ##############################
return $STATUS
```

E.4.5 Script: /usr/HACMP_ANSS/script/PRE_node_down_remote

```
#!/bin/ksh
# Program   : PRE_node_down_remote
# Role      : run before the event
# Arguments :   $1 = event name
#                  and the parameters passed in
# Written   : Wed Dec 13 16:43:24 CST 1995
# Modified  :
 . /usr/HACMP_ANSS/tools/tool_var
STATUS=0
(print †n=PRE-EVENT===============$(date)†
print †on : $(hostname) †
print †BEFORE : $1†
shift
print †Input Parameters: $*† )  >> $LOG
################################################################################
# Enter your customizing code here

######################## END OF CUSTOMIZATION ##############################
return $STATUS
```

E.4.6 Script: /usr/HACMP_ANSS/script/PRE_node_up_remote

```
#!/bin/ksh
# Program   : PRE_node_up_remote
# Role      : run before the event
# Arguments :   $1 = event name
#                  and the parameters passed in
# Written   : Wed Dec 13 16:50:41 CST 1995
# Modified  :
 . /usr/HACMP_ANSS/tools/tool_var
STATUS=0
(print †n=PRE-EVENT==============$(date)†
print †on : $(hostname) †
print †BEFORE : $1†
shift
print †Input Parameters: $*† )  >> $LOG
#############################################################################
# Enter your customizing code here

mail -s †Event Alert† thiess@thiessen.austin.ibm.com << END
Node goofy is about to re-enter the cluster. Users will be
migrated back from node mickey.
END

wall †Machine goofy has been recovered and is coming on-line.
There will be a short interruption for users of machine goofy.
Please logoff your application now.
You will be able to login to your application again within 5 minutes.†
sleep 10

###################### END OF CUSTOMIZATION ##########################
return $STATUS
```

E.4.7 Script: /usr/HACMP_ANSS/script/error_NOTIFICATION

```
#!/bin/ksh
#############################################################################
#
#   name             : error_NOTIFICATION
#   INPUT paremeters : $1 to $8 sent by errpt
#   Description      : called by each error, sends a message
#                                  into hacmp.errlog
#############################################################################
# Variables:
 . /usr/HACMP_ANSS/tools/tool_var
STATUS=0
G=$(tput smso)
F=$(tput rmso)
LOG=†$ERREURS/hacmp.errlog†
#####################################################################
#   main
#####################################################################
(print †************ Source and cause of error ****************†
print †HOSTNAME=$(hostname)  DATE=$(date)†
print †sequence number in error log = $1†
print †error ID                     = $2†
```

```
print †error class                  = $3†
print †error type                   = $4†
print †alert flag                   = $5†
print †resource name                = $6†
print †resource type                = $7†
print †resource class               = $8†
print †error label                  = $9†) >> $LOG
#############################################################################
# DO NOT FORGET TO set TO_WHOM in error_MAIL
 . /usr/HACMP_ANSS/tools/ERROR_TOOL/error_MAIL $1 $2 $3 $4 $5 $6 $7 $8 $9
#############################################################################
# DO NOT FORGET TO set QUEUE in error_PRINT
# . /usr/HACMP_ANSS/tools/ERROR_TOOL/error_PRINT $1 $2 $3 $4 $5 $6 $7 $8 $9
#############################################################################

return $STATUS
```

E.4.8 Script: /usr/HACMP_ANSS/script/error_SDA

```
#!/bin/ksh
#############################################################################################
# Written by: AUTOMATE
# Last modification by *** who ***
#
#  script:  error_SDA
#  parameters: 8 parameters (documented in error_NOTIFICATION)
#
# ARGUMENTS received :
# †sequence number in the error log = $1†
# †error ID                         = $2†
# †error class                      = $3†
# †error type                       = $4†
# †alert flag                       = $5†
# †resource name                    = $6†
# †resource type                    = $7†
# †resource class                   = $8†
# †error label                      = $9†
#############################################################################################
# Variables:
 . /usr/HACMP_ANSS/tools/tool_var
STATUS=0
( echo †n=error_SDA==============`date`†
echo †ERROR DETECTED: error_SDA†) |tee -a $ERREURS/hacmp.errlog> /dev/console
 . $SCRIPTS/error_NOTIFICATION
########################## START OF CUSTOMIZATION ###################################
#
LOCALNODENAME=$(/usr/sbin/cluster/utilities/get_local_nodename)
mail -s †Error Alert† thiess@thiessen.austin.ibm.com << END
An error has been detected on the HACMP cluster node $LOCALNODENAME
look at the $LOG file on the node.
DEVICE = $6
ADAPTER = $8

The system will be shut down and the users moved to a backup node.
END
```

```
wall †System will be shutting Down in 20 Seconds. Please log off now.
You will be able to login to your application again within 5 minutes.†
sleep 20

#  This command does a shutdown with takeover of HACMP

/usr/sbin/cluster/utilities/clstop -y ¢-N¢  ¢-gr¢
sleep 5

#  We now want to shutdown the machine, until our administrator can
#  investigate the problem.

/etc/shutdown -Fr

#########################  END OF CUSTOMIZATION  ###############################
return $STATUS
```

E.4.9 Script: /usr/HACMP_ANSS/script/event_NOTIFICATION

```
#!/bin/ksh
###############################################################################
#
#   name             : event_NOTIFICATION
#   INPUT paremeters : $1 = name of the event
#                      $2 = start or complete
#                      $3 = return code if $2 == complete
#                      all the arguments sent to the event
#
#   Description      : called by each event
###############################################################################
######### variables
 . /usr/HACMP_ANSS/tools/tool_var
STATUS=0
(print †n=NOTIFICATION==============$(date)†
print †on: $(hostname) † ) >> $LOG
if [ $2 = ¢start¢ ]
then
        quand=†START: $1†
    shift 2
    arguments=†arguments: $*†
else
        quand=†OUTPUT: $1 †
    arguments=†return code : $3†
fi
(print †$quand† ; print †$arguments† ) >> $LOG
###############################################################################
# DO NOT FORGET TO set TO_WHOM in event_MAIL
 . /usr/HACMP_ANSS/tools/EVENT_TOOL/event_MAIL $1 $2 $3
###############################################################################
# DO NOT FORGET TO set QUEUE in event_PRINT
#. /usr/HACMP_ANSS/tools/EVENT_TOOL/event_PRINT $1 $2 $3 $4 $5 $6 $7 $8
###############################################################################
return $STATUS
```

E.4.10 Script : /usr/HACMP_ANSS/tools/tool_var

```
HACMP=/usr/HACMP_ANSS
D=$HACMP/dessin
S=$HACMP/script
T=$HACMP/tools
U=$HACMP/utils
L=$HACMP/locks
G=$(tput smso)
N=$(tput rmso)
if [ ! -d $U ]
then mkdir $U
fi
#conf_var
##########################################

# Variables:  PRODUIT = directory containing HACMP commands
#             SCRIPTS = directory containing customized event scripts
#             ERREURS = directory where error messages are written
#             TOOLS   = directory containing the tools themselves
#             BACKUP  = directory where the original default scripts are saved
#             UTILS  = directory containing utilities used by the tools
PRODUIT=†/usr/sbin/cluster†
HACMP=†/usr/HACMP_ANSS†
SCRIPTS=†$HACMP/script†
TOOLS=†$HACMP/tools†
ERROR_TOOL=†$TOOLS/ERROR_TOOL†
EVENT_TOOL=†$TOOLS/EVENT_TOOL†
DOC_TOOL=†$TOOLS/DOC_TOOL†
CONF_TOOL=†$TOOLS/CONF_TOOL†
UTILS=†$HACMP/utils†
BACKUP=†$HACMP/backup†
DESSIN=†$HACMP/dessin†
LOCKS=†$HACMP/locks†
ERREURS=†/var/HACMP_ANSS/log†
if [ ! -d /usr/HACMP_ANSS/script ]
then
     mkdir /usr/HACMP_ANSS/script
fi
if [ ! -d /usr/HACMP_ANSS/backup ]
then
     mkdir /usr/HACMP_ANSS/backup
fi
if [ ! -d /usr/HACMP_ANSS/utils ]
then
     mkdir /usr/HACMP_ANSS/utils
fi
if [ ! -d /usr/HACMP_ANSS/locks ]
then
     mkdir /usr/HACMP_ANSS/locks
fi
export PATH=†$PATH:$TOOLS:$SCRIPTS:$PRODUIT:$UTILS†
LOG=†${ERREURS}/hacmp.eventlog†
```

E.5 SYSTEM FILES

E.5.1 File: /etc/rc

```
#!/bin/ksh
# @(#)06   1.13  com/cfg/etc/rc.sh, bos, bos320 4/30/91 14:25:11
#
# COMPONENT_NAME: (CFGETC) Multi-user mode system setup
#
# FUNCTIONS: rc
#
# ORIGINS: 27
#
# (C) COPYRIGHT International Business Machines Corp. 1989, 1990
# All Rights Reserved
# Licensed Materials - Property of IBM
#
# US Government Users Restricted Rights - Use, duplication or
# disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
#
#######################################################################

/usr/bin/dspmsg rc.cat 1 ¢Starting Multi-user Initializationn¢
PATH=/bin:/usr/bin:/usr/ucb:/etc::
ODMDIR=/etc/objrepos
export PATH ODMDIR

# Varyon all Volume Groups marked as auto-varyon.
# ( rootvg already varied on)
dspmsg rc.cat 2 ¢ Performing auto-varyon of Volume Groups n¢
/etc/cfgvg

# Activate all paging spaces in automatic list
# (those listed in /etc/swapspaces)
dspmsg rc.cat 3 ¢ Activating all paging spaces n¢
/etc/swapon -a

# Perform file system checks
# The -f flag skips the check if the log has been replayed successfully
fsck -fp

# Perform all auto mounts
dspmsg rc.cat 4 ¢ Performing all automatic mounts n¢
mount all

# Remove /etc/nologin if left behind by shutdown
rm -f /etc/nologin

# Running expreserve to recover vi editor sessions
/usr/lib/expreserve - 2>/dev/null

# Write a dummy record to file /usr/adm/sa/sa<date> to specify
# that system start up has occurred.
```

```
# dspmsg rc.cat 6 ¢Write system start up record to /usr/adm/sa/sa`date`¢
#/bin/su - root -c /usr/lib/sa/sadc /usr/adm/sa/sa`date +%d`

# Manufacturing post install process.
# This must be at the end of this file, /etc/rc.
if [ -x /etc/mfg/rc.preload ]
then
     /etc/mfg/rc.preload
fi

dspmsg rc.cat 5 ¢Multi-user initialization completedn¢
exit 0
```

E.5.2 File: /etc/rc.net

```
#!/bin/ksh
# @(#)90        1.18  com/cmd/net/netstart/rc.net, cmdnet, bos320, 9150320k 12/11/91 14:40
:04
#
# COMPONENT_NAME: CMDNET   (/etc/rc.net)
#
# ORIGINS: 27
#
# (C) COPYRIGHT International Business Machines Corp. 1985, 1989
# All Rights Reserved
# Licensed Materials - Property of IBM
#
# US Government Users Restricted Rights - Use, duplication or
# disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
#
# HACMP6000
# HACMP6000 These lines added by HACMP6000 software
[ †$1† = †-boot† ] && shift || exit 0 # HACMP6000
# HACMP6000

#########################################################################
# rc.net - called by cfgmgr during 2nd boot phase.
#
# Configures and starts TCP/IP interfaces.
# Sets hostname, default gateway and static routes.
# Note: all the stdout should be redirected to a file (e.g. /dev/null),
#    because stdout is used to pass logical name(s) back to the cfgmgr
#    to be configured.  The LOGFILE variable specifies the output file.
# The first section of rc.net configures the network via the new
#    configuration methods.  These configuration methods require that
#    the interface and protocol information be entered in the ODM
#    database (with either SMIT or the high level configuration commands
#    (mkdev, chdev).
# The second section (commented out) is an example of the equivalent
#    traditional commands used to perform the same function.  You may
#    use the traditional commands instead of the configuration methods
#    if you prefer.  These commands do NOT use the ODM database.
# The third section performs miscellaneous commands which are
#    compatible with either of the previous two sections.
#########################################################################
```

```
#
# Close file descriptor 1 and 2 because the parent may be waiting
# for the file desc. 1 and 2 to be closed.  The reason is that this shell
# script may spawn a child which inherit all the file descriptor from the parent
# and the child process may still be running after this process is terminated.
# The file desc. 1 and 2 are not closed and leave the parent hanging
# waiting for those desc. to be finished.
#LOGFILE=/dev/null   # LOGFILE is where all stdout goes.
LOGFILE=/tmp/rc.net.out    # LOGFILE is where all stdout goes.
>$LOGFILE       # truncate LOGFILE.
exec 1<&-       # close descriptor 1
exec 2<&-       # close descriptor 2
exec 1< /dev/null    # open descriptor 1
exec 2<    /dev/null  # open descriptor 2

no -d lowclust       # set cluster low water mark


#####################################################################
# Part I - Configuration using the data in the ODM database:
# Enable network interface(s):
#####################################################################
# This should be done before routes are defined.
# For each network adapter that has already been configured, the
# following commands will define, load and configure a corresponding
# interface.
/usr/lib/methods/defif               >>$LOGFILE 2>&1
/usr/lib/methods/cfgif  $*                >>$LOGFILE 2>&1

#####################################################################
#  Special X25 and SLIP handling
#####################################################################
# In addition to configure the network interface,  X25 and SLIP
# interfaces require special commands to complete the configuration
# The x25xlate command bring the x25 translation table into the
# kernel while the slattach changes the tty handling for the tty
# port used by the the SLIP interface.   A separate slattach command is
# execute for every tty port used by configured SLIP interfaces.

X25HOST=`lsdev -C -c if -s XT -t xt -S available`
if [ ! -z "$X25HOST" ]
then
     x25xlate             >>$LOGFILE 2>&1
fi

SLIPHOST=`lsdev -C -c if -s SL -t sl -S available | awk '{ print $1 }'`
for i in $SLIPHOST
do
     echo $i              >>$LOGFILE 2>&1
     TTYPORT=`lsattr -E -l $i -F value -a ttyport`
     TTYBAUD=`lsattr -E -l $i -F value -a baudrate`
     TTYDIALSTRING=`lsattr -E -l $i -F value -a dialstring`
     rm -f /etc/locks/LCK..$TTYPORT
     if [ -z "$TTYBAUD" -a  -z "$TTYDIALSTRING" ]
     then
```

```
          FromHOST=`lsattr -E -l $i -F tvaluet -a netaddr`
          DestHOST=`lsattr -E -l $i -F tvaluet -a dest`
          SLIPMASK=`lsattr -E -l $i -F tvaluet -a netmask`
          if [ -z t$SLIPMASKt ]
          then
               ifconfig $SLIPHOST inet $FromHOST $DestHOST up
          else
               ifconfig $SLIPHOST inet $FromHOST $DestHOST netmask $SLIPMASK up
          fi
          ( slattach $TTYPORT )                >>$LOGFILE 2>&1
     else
          eval DST=¢$TTYDIALSTRING¢           >>$LOGFILE 2>&1
          ( eval slattach $TTYPORT $TTYBAUD ¢$DST¢ ) >>$LOGFILE 2>>$LOGFILE
     fi
done

#########################################################################
# Configure the Internet protocol kernel extension (netinet):
#########################################################################
# The following commands will also set hostname, default gateway,
# and static routes as found in the ODM database for the network.
/usr/lib/methods/definet             >>$LOGFILE 2>&1
/usr/lib/methods/cfginet             >>$LOGFILE 2>&1


#########################################################################
# Part II - Traditional Configuration.
#########################################################################
# An alternative method for bringing up all the default interfaces
# is to specify explicitly which interfaces to configure using the
# ifconfig command.  Ifconfig requires the configuration information
# be specified on the command line.  Ifconfig will not update the
# information kept in the ODM configuration database.
#
# Valid network interfaces are:
# lo=local loopback, en=standard ethernet, et=802.3 ethernet
# sl=serial line IP, tr=802.5 token ring, xt=X.25
#
# e.g., en0 denotes standard ethernet network interface, unit zero.
#
# Below are examples of how you could bring up each interface using
# ifconfig.  Since you can specify either a hostname or a dotted
# decimal address to set the interface address, it is convenient to
# set the hostname at this point and use it for the address of
# an interface, as shown below:
#
#/bin/hostname robo.austin.ibm.com   >>$LOGFILE 2>&1
#
# (Remember that if you have more than one interface,
# you¢ll want to have a different IP address for each one.
# Below, xx.xx.xx.xx stands for the internet address for the
# given interface).
#
#/usr/sbin/ifconfig lo0 inet loopback    up >>$LOGFILE 2>&1
#/usr/sbin/ifconfig en0 inet `hostname`  up >>$LOGFILE 2>&1
```

```
#/usr/sbin/ifconfig et0 inet xx.xx.xx.xx  up >>$LOGFILE 2>&1
#/usr/sbin/ifconfig tr0 inet xx.xx.xx.xx  up >>$LOGFILE 2>&1
#/usr/sbin/ifconfig sl0 inet xx.xx.xx.xx  up >>$LOGFILE 2>&1
#/usr/sbin/ifconfig xt0 inet xx.xx.xx.xx  up >>$LOGFILE 2>&1
#
#
# Now we set any static routes.
#
# /usr/sbin/route add 0 gateway            >>$LOGFILE 2>&1
# /usr/sbin/route add 192.9.201.0 gateway       >>$LOGFILE 2>&1


#####################################################################
# Part III - Miscellaneous Commands.
#####################################################################
# Set the hostid and uname to `hostname`, where hostname has been
# set via ODM in Part I, or directly in Part II.
# (Note it is not required that hostname, hostid and uname all be
# the same).
/usr/sbin/hostid `hostname`          >>$LOGFILE 2>&1
/bin/uname -S`hostname|sed ¢s/..*$//¢`    >>$LOGFILE 2>&1

#####################################################
# The socket default buffer size (initial advertized TCP window) is being
# set to a default value of 16k (16384). This improves the performance
# for ethernet and token ring networks.  Networks with lower bandwidth
# such as SLIP (Serial Line Internet Protocol) and X.25 or higher bandwidth
# such as Serial Optical Link and FDDI would have a different optimum
# buffer size.
# ( OPTIMUM WINDOW = Bandwidth * Round Trip Time )
#####################################################
if [ -f /usr/sbin/no ] ; then
     /usr/sbin/no -o tcp_sendspace=16384
     /usr/sbin/no -o tcp_recvspace=16384
fi

/etc/no -o ipforwarding=0
/etc/no -o ipsendredirects=0
```

E.5.3 File: /etc/hosts

```
# @(#)47   1.1  com/cmd/net/netstart/hosts, bos, bos320 7/24/91 10:00:46
#
# COMPONENT_NAME: TCPIP hosts
#
# FUNCTIONS: loopback
#
# ORIGINS: 26  27
#
# (C) COPYRIGHT International Business Machines Corp. 1985, 1989
# All Rights Reserved
# Licensed Materials - Property of IBM
#
# US Government Users Restricted Rights - Use, duplication or
# disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
```

```
#
#  /etc/hosts
#
# This file contains the hostnames and their address for hosts in the
# network.  This file is used to resolve a hostname into an Internet
# address.
#
# At minimum, this file must contain the name and address for each
# device defined for TCP in your /etc/net file.  It may also contain
# entries for well-known (reserved) names such as timeserver
# and printserver as well as any other host name and address.
#
# The format of this file is:
# Internet Address   Hostname   # Comments
# Items are separated by any number of blanks and/or tabs.  A ¢#¢
# indicates the beginning of a comment; characters up to the end of the
# line are not interpreted by routines which search this file.  Blank
# lines are allowed.

# Internet Address   Hostname   # Comments
# 192.9.200.1        net0sample      # ethernet name/address
# 128.100.0.1        token0sample    # token ring name/address
# 10.2.0.2           x25sample  # x.25 name/address
127.0.0.1       loopback localhost   # loopback (lo0) name/address

# Cluster 1 - disney

9.3.1.79   mickey.itsc.austin.ibm.com mickey
9.3.4.79   mickey_sb.itsc.austin.ibm.com mickey_sb
9.3.5.79   mickey_en.itsc.austin.ibm.com mickey_en

9.3.1.46   goofy_boot.itsc.austin.ibm.com goofy_boot
9.3.1.80   goofy.itsc.austin.ibm.com goofy
9.3.4.80   goofy_sb.itsc.austin.ibm.com goofy_sb
9.3.5.80   goofy_en.itsc.austin.ibm.com goofy_en

# Cluster 2 - dave

9.3.1.3         hadave1_boot.itsc.austin.ibm.com hadave1_boot
9.3.1.16   hadave1.itsc.austin.ibm.com hadave1
9.3.4.16    hadave1_sb.itsc.austin.ibm.com hadave1_sb

9.3.1.6         hadave2_boot.itsc.austin.ibm.com hadave2_boot
9.3.1.17   hadave2.itsc.austin.ibm.com hadave2
9.3.4.17   hadave2_sb.itsc.austin.ibm.com hadave2_sb

# Client & Others

9.3.1.43        pluto
9.3.1.74        gandalf
9.209.46.194    surveyor
9.209.41.111    aix11
9.209.32.4           jd560
9.3.4.16   hadave1_sb.itsc.austin.ibm.com hadave1_sb
```

```
9.3.1.3    hadave1_boot.itsc.austin.ibm.com hadave1_boot
9.3.1.45   mickey_boot.itsc.austin.ibm.com mickey
```

E.5.4 File: /etc/filesystems

```
* @(#)filesystems @(#)29   1.18  com/cfg/etc/filesystems, bos, bos320 8/21/91 08:32:3
1
*
* COMPONENT_NAME: CFGETC
*
* FUNCTIONS:
*
* ORIGINS: 27
*
* (C) COPYRIGHT International Business Machines Corp. 1985, 1991
* All Rights Reserved
* Licensed Materials - Property of IBM
*
* US Government Users Restricted Rights - Use, duplication or
* disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
*
*
*
* This version of /etc/filesystems assumes that only the root file system
* is created and ready.  As new file systems are added, change the check,
* mount, free, log, vol and vfs entries for the appropriate stanza.
*

/:
     dev        = /dev/hd4
     vfs        = jfs
     log        = /dev/hd8
     mount           = automatic
     check           = false
     type       = bootfs
     vol        = root
     free       = true

/home:
     dev       = /dev/hd1
     vol       = †/home†
     mount     = true
     check     = true
     free      = false
     vfs    = jfs
     log    = /dev/hd8

/usr:
     dev        = /dev/hd2
     vfs        = jfs
     log        = /dev/hd8
     mount           = automatic
     check           = false
     type       = bootfs
     vol        = /usr
```

```
     free      = false

/var:
     dev       = /dev/hd9var
     vol       = †/var†
     mount     = automatic
     check     = false
     free      = false
     vfs    = jfs
     log    = /dev/hd8
     type   = bootfs

/tmp:
     dev       = /dev/hd3
     vfs       = jfs
     log       = /dev/hd8
     mount          = automatic
     check          = false
     vol       = /tmp
     free      = false

/mnt:
     dev  = /dev/hd7
     vol  = †spare†
     mount     = false
     check     = false
     free = false
     vfs  = jfs
     log     = /dev/hd8

/blv:
     dev  = /dev/hd5
     vol  = †spare†
     mount     = false
     check     = false
     free = false
     vfs  = jfs
     log     = /dev/hd8

/usr/bin/blv.fs:
     dev  = /usr/bin/blv.fs
     vol  = †/†

/inst:
     dev       = /dev/extlv1
     vfs       = jfs
     log       = /dev/extloglv
     mount          = false
     check          = false
     options        = rw
     account        = false

/test1:
     dev       = /dev/lvtest1
     vfs       = jfs
```

```
     log       = /dev/loglvtest1
     mount          = false
     check          = false
     options        = rw
     account        = false

/test2:
     dev       = /dev/lvtest2
     vfs       = jfs
     log       = /dev/loglvtest2
     mount          = false
     check          = false
     options        = rw
     account        = false
```

E.5.5 File: /etc/inetd.conf

```
#
# COMPONENT_NAME: TCPIP inetd.conf
#
# FUNCTIONS:
#
# ORIGINS: 26  27
#
# (C) COPYRIGHT International Business Machines Corp. 1985, 1989
# All Rights Reserved
# Licensed Materials - Property of IBM
#
# US Government Users Restricted Rights - Use, duplication or
# disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
#
#  /etc/inetd.conf
#
#                 Internet server configuration database
#
#    Services can be added and deleted by deleting or inserting a
#    comment character (ie. #) at the beginning of a line  If inetd
#    is running under SRC control then the tinetimpt command must
#    be executed to import the information from this file to the
#    InetServ ODM object class, then the trefresh -s inetdt command
#    needs to be executed for inetd to re-read the InetServ database.
#
#    NOTE: The TCP/IP servers do not require SRC and may be started
#    by invoking the service directly (i.e. /etc/inetd). If inetd
#    has been invoked directly, after modifying this file, send a
#    hangup signal, SIGHUP to inetd (ie. kill -1 tpid_of_inetdt).
#
#    require that the portmap daemon be running.
#
# service   socket   protocol   wait/   user    server    server program
#  name      type              nowait          program     arguments
#
## The following line is the new style tftp daemon - allows write create.
## The following line needs to be uncommented and run inetimp to enable tftpd
## The following line is for installing over the network.
```

```
echo stream     tcp  nowait    root internal
echo dgram      udp  wait root internal
discard    stream     tcp  nowait    root internal
discard    dgram      udp  wait root internal
daytime    stream     tcp  nowait    root internal
daytime    dgram      udp  wait root internal
chargen    stream     tcp  nowait    root internal
chargen    dgram      udp  wait root internal
ftp  stream     tcp  nowait    root /etc/ftpd ftpd
telnet     stream     tcp  nowait    root /etc/telnetd telnetd
time stream     tcp  nowait    root internal
time dgram      udp  wait root internal
#bootps    dgram      udp  wait root /etc/bootpd bootpd
#tftp      dgram      udp  wait nobody    /etc/tftpd tftpd -n
#finger    stream     tcp  nowait    nobody    /etc/fingerd fingerd
#rexd      sunrpc_tcp tcp  wait root /usr/etc/rpc.rexd rexd 100017 1
executiond      sunrpc_tcp tcp  wait root /usr/lpp/sd/executiond executiond 300201 1

comp_ed    sunrpc_tcp tcp  wait root /usr/lpp/sd/executiond comp_ed 33333332 1
rstatd     sunrpc_udp udp  wait root /usr/etc/rpc.rstatd rstatd 100001 1-3
rusersd    sunrpc_udp udp  wait root /usr/etc/rpc.rusersd rusersd 100002 1-2
rwalld     sunrpc_udp udp  wait root /usr/etc/rpc.rwalld rwalld 100008 1
sprayd     sunrpc_udp udp  wait root /usr/etc/rpc.sprayd sprayd 100012 1
pcnfsd     sunrpc_udp udp  wait root /etc/rpc.pcnfsd pcnfsd 150001 1
exec stream     tcp  nowait    root /etc/rexecd rexecd
#biff      dgram      udp  wait root /etc/comsat comsat
login      stream     tcp  nowait    root /etc/rlogind rlogind
shell      stream     tcp  nowait    root /etc/rshd rshd
#talk      dgram      udp  wait root /etc/talkd talkd
ntalk      dgram      udp  wait root /etc/talkd talkd
uucp stream     tcp  nowait    root /etc/uucpd uucpd
#instsrv   stream     tcp  nowait    netinst   /u/netinst/bin/instsrv instsrv -r
/tmp/netinstalllog /u/netinst/scripts
godm stream     tcp  nowait    root /usr/sbin/cluster/godmd
```

E.5.6 File: /etc/syslog.conf

```
# @(#)34   1.9  com/cmd/net/syslogd/syslog.conf, cmdnet, bos325, 9331325b 6/13/93 14:
52:39
#
# COMPONENT_NAME: (CMDNET) Network commands.
#
# FUNCTIONS:
#
# ORIGINS: 27
#
# (C) COPYRIGHT International Business Machines Corp. 1988, 1989
# All Rights Reserved
# Licensed Materials - Property of IBM
#
# US Government Users Restricted Rights - Use, duplication or
# disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
#
# /etc/syslog.conf - control output of syslogd
#
```

```
#
# Each line must consist of two parts:-
#
# 1) A selector to determine the message priorities to which the
#    line applies
# 2) An action.
#
# The two fields must be separated by one or more tabs or spaces.
#
# format:
#
# <msg_src_list>          <destination>
#
# where <msg_src_list> is a semicolon separated list of <facility>.<priority>
# where:
#
# <facility> is:
#    * - all (except mark)
#    mark - time marks
#    kern,user,mail,daemon, auth,... (see syslogd(AIX Commands Reference))
#
# <priority> is one of (from high to low):
#    emerg/panic,alert,crit,err(or),warn(ing),notice,info,debug
#    (meaning all messages of this priority or higher)
#
# <destination> is:
#    /filename - log to this file
#    username[,username2...] - write to user(s)
#    @hostname - send to syslogd on this machine
#    * - send to all logged in users
#
# example:
# †mail messages, at debug or higher, go to Log file. File must exist.†
# †all facilities, at debug and higher, go to console†
# †all facilities, at crit or higher, go to all users†
#  mail.debug        /usr/spool/mqueue/syslog
#  *.debug           /dev/console
#  *.crit            *
# HACMP/6000 Critical Messages from HACMP/6000
local0.crit /dev/console
# HACMP/6000 Informational Messages from HACMP/6000
local0.info /usr/adm/cluster.log
# HACMP/6000 Messages from Cluster Scripts
user.notice /usr/adm/cluster.log
```

E.5.7 File: /etc/inittab

```
: @(#)49  1.28  com/cfg/etc/inittab, bos, bos320 10/3/91 10:46:51
:  COMPONENT_NAME: CFGETC
:
:  ORIGINS: 3, 27
:
:  (C) COPYRIGHT International Business Machines Corp. 1989, 1990
:  All Rights Reserved
:  Licensed Materials - Property of IBM
```

```
:
:  US Government Users Restricted Rights - Use, duplication or
:  disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
:
: Note - initdefault and sysinit should be the first and second entry.
:
init:2:initdefault:
brc::sysinit:/sbin/rc.boot 3 >/dev/console 2>&1 # Phase 3 of system boot
powerfail::powerfail:/etc/rc.powerfail >/dev/console 2>&1  # d51225
rc:2:wait:/etc/rc > /dev/console 2>&1  # Multi-User checks
fbcheck:2:wait:/usr/lib/dwm/fbcheck >/dev/console 2>&1 # run /etc/firstboot
srcmstr:2:respawn:/etc/srcmstr       # System Resource Controller
harc:2:wait:/usr/sbin/cluster/etc/harc.net # HACMP6000 network startup
rctcpip: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
cons:0123456789:respawn:/etc/getty /dev/console
piobe:2:wait:/bin/rm -f /usr/lpd/pio/flags/*  # Clean up printer flags files
cron:2:respawn:/etc/cron
qdaemon:a:wait:/bin/startsrc -sqdaemon
writesrv:a:wait:/bin/startsrc -swritesrv
uprintfd:2:respawn:/etc/uprintfd
rcncs:a:wait:sh /etc/rc.ncs
infod:2:once:startsrc -s infod
tty0:2:off:/etc/getty /dev/tty0
clvm6000:2:wait:/usr/sbin/cluster/cllvm -c status # Check CLVM stat
clinit:a:wait:touch /usr/sbin/cluster/.telinit # HACMP6000 This must be last entry in inittab!
```

E.6 CONTENTS OF THE HACMP OBJECTS IN THE ODM

E.6.1 odmget of /etc/objrepos/HACMPadapter

```
HACMPadapter:
      type = †ether†
      network = †etnet1†
      nodename = †goofy†
      ip_label = †goofy_ent†
      function = †service†
      identifier = †9.3.5.80†
      haddr = †† 

HACMPadapter:
      type = †rs232†
      network = †rsnet1†
      nodename = †goofy†
      ip_label = †goofy_tty0†
      function = †service†
      identifier = †/dev/tty0†
      haddr = †† 

HACMPadapter:
      type = †token†
      network = †trnet1†
      nodename = †goofy†
      ip_label = †goofy†
      function = †service†
      identifier = †9.3.1.80†
      haddr = †0x42005aa8d1f3† 

HACMPadapter:
      type = †token†
      network = †trnet1†
      nodename = †goofy†
      ip_label = †goofy_boot†
      function = †boot†
      identifier = †9.3.1.46†
      haddr = †† 

HACMPadapter:
      type = †token†
      network = †trnet1†
      nodename = †goofy†
      ip_label = †goofy_sb†
      function = †standby†
      identifier = †9.3.4.80†
      haddr = †† 

HACMPadapter:
      type = †ether†
      network = †etnet1†
      nodename = †mickey†
      ip_label = †mickey_ent†
```

```
      function = †service†
      identifier = †9.3.5.79†
      haddr = ††

HACMPadapter:
      type = †rs232†
      network = †rsnet1†
      nodename = †mickey†
      ip_label = †mickey_tty0†
      function = †service†
      identifier = †/dev/tty0†
      haddr = ††

HACMPadapter:
      type = †token†
      network = †trnet1†
      nodename = †mickey†
      ip_label = †mickey†
      function = †service†
      identifier = †9.3.1.79†
      haddr = †0x42005aa8b484†

HACMPadapter:
      type = †token†
      network = †trnet1†
      nodename = †mickey†
      ip_label = †mickey_boot†
      function = †boot†
      identifier = †9.3.1.45†
      haddr = ††

HACMPadapter:
      type = †token†
      network = †trnet1†
      nodename = †mickey†
      ip_label = †mickey_sb†
      function = †standby†
      identifier = †9.3.4.79†
      haddr = ††
```

E.6.2 odmget of /etc/objrepos/HACMPcluster

```
HACMPcluster:
      id = 1
      name = †disney†
      nodename = †mickey†
```

E.6.3 odmget of /etc/objrepos/HACMPcommand

```
HACMPcommand:
      command = †clverify†
      options = †software†
      optflag = 1
      path = ††
```

```
      numargs = 0
      args = ††
      help = †Tools for verifying that a cluster is properly installed and configured†
      catalog = †command.cat†
      setno = 0
      msgno = 2

HACMPcommand:
      command = †clverify†
      options = †cluster†
      optflag = 1
      path = ††
      numargs = 0
      args = ††
      help = †Tools for verifying that a cluster is properly installed and configured†
      catalog = †command.cat†
      setno = 0
      msgno = 3

HACMPcommand:
      command = †clverify.software†
      options = †bos†
      optflag = 1
      path = ††
      numargs = 0
      args = ††
      help = †Verifies that your software environment is compatible with HACMP†
      catalog = †command.cat†
      setno = 0
      msgno = 6

HACMPcommand:
      command = †clverify.software†
      options = †prereq†
      optflag = 1
      path = ††
      numargs = 0
      args = ††
      help = †Verifies that your software environment is compatible with HACMP†
      catalog = †command.cat†
      setno = 0
      msgno = 7

HACMPcommand:
      command = †clverify.software†
      options = †badptfs†
      optflag = 1
      path = ††
      numargs = 0
      args = ††
      help = †Verifies that your software environment is compatible with HACMP†
      catalog = †command.cat†
      setno = 0
      msgno = 8
```

```
HACMPcommand:
      command = †clverify.softwaret
      options = †lppt
      optflag = 1
      path = tt
      numargs = 0
      args = tt
      help = †Verifies that your software environment is compatible with HACMPt
      catalog = tcommand.catt
      setno = 0
      msgno = 8

HACMPcommand:
      command = †clverify.clustert
      options = ttopologyt
      optflag = 1
      path = tt
      numargs = 0
      args = tt
      help = †Verifies that your cluster is configured properlyt
      catalog = tcommand.catt
      setno = 0
      msgno = 9

HACMPcommand:
      command = †clverify.clustert
      options = tconfigt
      optflag = 1
      path = tt
      numargs = 0
      args = tt
      help = †Verifies that your cluster is configured properlyt
      catalog = tcommand.catt
      setno = 0
      msgno = 10

HACMPcommand:
      command = †clverify.software.prereqt
      options = tt
      optflag = 0
      path = t/usr/sbin/cluster/diag/clvreqt
      numargs = 0
      args = tt
      help = †Verifies that all fixes to AIX required by HACMP have been installedt
      catalog = tcommand.catt
      setno = 0
      msgno = 13

HACMPcommand:
      command = †clverify.software.lppt
      options = tt
      optflag = 0
      path = t/usr/sbin/cluster/diag/clvhacmpt
      numargs = 0
      args = tt
```

```
        help = †Verifies that HACMP is properly installed†
        catalog = †command.cat†
        setno = 0
        msgno = 14

HACMPcommand:
        command = †clverify.software.bos†
        options = ††
        optflag = 0
        path = †/usr/sbin/cluster/diag/clvbos†
        numargs = 0
        args = ††
        help = †Verifies that the AIX level is correct for HACMP†
        catalog = †command.cat†
        setno = 0
        msgno = 15

HACMPcommand:
        command = †clverify.software.badptfs†
        options = ††
        optflag = 0
        path = †/usr/sbin/cluster/diag/clvinval†
        numargs = 0
        args = ††
        help = †Verifies that no known PTFs that break HACMP are installed†
        catalog = †command.cat†
        setno = 0
        msgno = 16

HACMPcommand:
        command = †clverify.cluster.topology†
        options = †check†
        optflag = 1
        path = ††
        numargs = 0
        args = ††
        help = †Verifies that all cluster nodes agree on cluster topology†
        catalog = †command.cat†
        setno = 0
        msgno = 17

HACMPcommand:
        command = †clverify.cluster.topology†
        options = †sync†
        optflag = 1
        path = ††
        numargs = 0
        args = ††
        help = †Forces all cluster nodes to agree on cluster topology†
        catalog = †command.cat†
        setno = 0
        msgno = 18

HACMPcommand:
        command = †clverify.cluster.topology.check†
```

```
      options = †† 
      optflag = 0
      path = †/usr/sbin/cluster/diag/clconfig†
      numargs = 1
      args = †-t†
      help = †Verifies that all cluster nodes agree on cluster topology†
      catalog = †command.cat†
      setno = 0
      msgno = 19

HACMPcommand:
      command = †clverify.cluster.topology.sync†
      options = †† 
      optflag = 0
      path = †/usr/sbin/cluster/diag/clconfig†
      numargs = 2
      args = †-s -t†
      help = †Forces all cluster nodes to agree on cluster topology†
      catalog = †command.cat†
      setno = 0
      msgno = 20

HACMPcommand:
      command = †clverify.cluster.config†
      options = †networks†
      optflag = 1
      path = †† 
      numargs = 0
      args = †command.cat†
      help = †Verifies that cluster resources are properly installed†
      catalog = †† 
      setno = 0
      msgno = 23

HACMPcommand:
      command = †clverify.cluster.config†
      options = †resources†
      optflag = 1
      path = †† 
      numargs = 0
      args = †† 
      help = †Verifies that cluster resources are properly installed†
      catalog = †command.cat†
      setno = 0
      msgno = 22

HACMPcommand:
      command = †clverify.cluster.config†
      options = †both†
      optflag = 1
      path = †† 
      numargs = 0
      args = †† 
      help = †Verifies that cluster resources are properly installed†
      catalog = †command.cat†
```

```
      setno = 0
      msgno = 21

HACMPcommand:
      command = †clverify.cluster.config.networks†
      options = ††
      optflag = 0
      path = †/usr/sbin/cluster/diag/clconfig†
      numargs = 2
      args = †-v -t†
      help = †Checks for proper configuration of network adapters and tty lines†
      catalog = †command.cat†
      setno = 0
      msgno = 25

HACMPcommand:
      command = †clverify.cluster.config.resources†
      options = ††
      optflag = 0
      path = †/usr/sbin/cluster/diag/clconfig†
      numargs = 2
      args = †-v -r†
      help = †Checks for agreement on resource ownership and takeover distribution†
      catalog = †command.cat†
      setno = 0
      msgno = 26

HACMPcommand:
      command = †clverify.cluster.config.both†
      options = ††
      optflag = 0
      path = †/usr/sbin/cluster/diag/clconfig†
      numargs = 1
      args = †-v†
      help = †Runs both the networks and resources programs†
      catalog = †command.cat†
      setno = 0
      msgno = 24

HACMPcommand:
      command = †cldiag†
      options = †logs†
      optflag = 1
      path = ††
      numargs = 0
      args = ††
      help = †Allows for selected viewing of HACMP log files, enables debugging of the C
luster Manager, or enables dumping of all Lock Manager resources.†
      catalog = †command.cat†
      setno = 0
      msgno = 27

HACMPcommand:
      command = †cldiag.logs†
      options = †scripts†
```

```
      optflag = 1
      path = †
      numargs = 0
      args = †
      help = †Allows for selected viewing of script output or syslog output.†
      catalog = †command.cat†
      setno = 0
      msgno = 28

HACMPcommand:
      command = †cldiag.logs.scripts†
      options = †
      optflag = 0
      path = †/usr/sbin/cluster/diag/cld_logfiles†
      numargs = 2
      args = †-t scripts†
      help = †scripts [-h host] [-s] [-f] [-d days] [-R file] [event ...]
where:
      -h host  is the name of a remote host from which to gather log data
      -s       filters Start/Complete events
      -f       filters failure events
      -d days  defines the number of previous days from which to retrieve log
      -R file  is file to which output is saved
      event     is a list of cluster events

Allows for parsing the /tmp/hacmp.out file
†
      catalog = †command.cat†
      setno = 0
      msgno = 29

HACMPcommand:
      command = †cldiag.logs†
      options = †syslog†
      optflag = 1
      path = †
      numargs = 0
      args = †
      help = †Allows for selected viewing of script output or syslog output.†
      catalog = †command.cat†
      setno = 0
      msgno = 30

HACMPcommand:
      command = †cldiag.logs.syslog†
      options = †
      optflag = 0
      path = †/usr/sbin/cluster/diag/cld_logfiles†
      numargs = 2
      args = †-t syslog†
      help = †syslog [-h host] [-e] [-w] [-d days] [-R file] [process ...]
where:
      -h host  is the name of a remote host from which to gather log data
      -e       filters error events
      -w       filters warning events
```

```
      -d days  defines the number of previous days from which to retrieve log
      -R file  is file to which output is saved
      process  is a list of cluster daemon processes

Allows for parsing the /usr/adm/cluster.log file.
†
      catalog = †command.cat†
      setno = 0
      msgno = 31

HACMPcommand:
      command = †cldiag†
      options = †debug†
      optflag = 1
      path = ††
      numargs = 0
      args = ††
      help = †Allows for selected viewing of HACMP log files, enables debugging of the C
luster Manager, or enables dumping of all Lock Manager resources.†
      catalog = †command.cat†
      setno = 0
      msgno = 32

HACMPcommand:
      command = †cldiag.debug†
      options = †clstrmgr†
      optflag = 1
      path = ††
      numargs = 0
      args = ††
      help = †Enables debugging of the Cluster Manager or the dumping of the lock resour
ce table.†
      catalog = †command.cat†
      setno = 0
      msgno = 33

HACMPcommand:
      command = †cldiag.debug.clstrmgr†
      options = ††
      optflag = 0
      path = †/usr/sbin/cluster/diag/cld_debug†
      numargs = 2
      args = †-t clstrmgr†
      help = †clstrmgr [-l level] [-R file]
where:
      -l level  is the level of debugging performed (0 - 9, where 0 turns debugging off)

      -R file   is the file to which output is saved

Allows for real-time clstrmgr debugging.
†
      catalog = †command.cat†
      setno = 0
      msgno = 34
```

```
HACMPcommand:
      command = †cldiag.debug†
      options = †cllockd†
      optflag = 1
      path = ††
      numargs = 0
      args = ††
      help = †Enables debugging of the Cluster Manager or the dumping of the lock resour
ce table.†
      catalog = †command.cat†
      setno = 0
      msgno = 35

HACMPcommand:
      command = †cldiag.debug.cllockd†
      options = ††
      optflag = 0
      path = †/usr/sbin/cluster/diag/cld_debug†
      numargs = 2
      args = †-t cllockd†
      help = †cllockd [-R file]
where:
      -R file   is the file to which output is saved

Allows dumping of the Lock Resource Table.
†
      catalog = †command.cat†
      setno = 0
      msgno = 36

HACMPcommand:
      command = †cldiag†
      options = †vgs†
      optflag = 1
      path = ††
      numargs = 0
      args = ††
      help = †Finds volume group inconsistencies among hosts and the disks.†
      catalog = †command.cat†
      setno = 0
      msgno = 37

HACMPcommand:
      command = †cldiag.vgs†
      options = ††
      optflag = 0
      path = †/usr/sbin/cluster/diag/cld_vgs†
      numargs = 0
      args = ††
      help = †vgs hostnames [-v volume_groups]
where:
      -h hostnames      is a list of 2 to 8 hostnames separated by commas
      -v volume_groups  is a list of volume group names separated by commas
      Note: Spaces are not allowed between hostname entries or volume group entries
```

```
Checks for consistencies of volume groups among hosts, ODMs, and disks.
†
      catalog = †command.cat†
      setno = 0
      msgno = 38

HACMPcommand:
      command = †cldiag†
      options = †trace†
      optflag = 1
      path = ††
      numargs = 0
      args = ††
      help = †Obtains a sequential flow of time stamped system events.†
      catalog = †command.cat†
      setno = 0
      msgno = 39

HACMPcommand:
      command = †cldiag.trace†
      options = ††
      optflag = 0
      path = †/usr/sbin/cluster/diag/cld_trace†
      numargs = 0
      args = ††
      help = †trace [-t time] [-R file] [-l] daemon ...
where:
      -t time  is the number of seconds to perform the trace
      -R file  is file to which output is saved
      -l       chooses a more detailed trace option
      daemon   is a list of cluster daemons to trace

Allows for tracing HACMP daemons (clstrmgr, cllockd, clsmuxpd, clinfo).
†
      catalog = †command.cat†
      setno = 0
      msgno = 40

HACMPcommand:
      command = †cldiag†
      options = †error†
      optflag = 1
      path = ††
      numargs = 0
      args = ††
      help = †Displays errors from the error log (hardware, software, system) that occur
 in the cluster.†
      catalog = †command.cat†
      setno = 0
      msgno = 41

HACMPcommand:
      command = †cldiag.error†
      options = ††
      optflag = 0
```

```
      path = †/usr/sbin/cluster/diag/cld_error†
      numargs = 0
      args = ††
      help = †error type [-h host] [-R file]
where:
      type is one of:
            short    - short eror report
            long     - long error report
            cluster  - HACMP/6000 specific short error report
      -h host  is the name of a remote host from which to gather log data
      -R file  is file to which output is saved

Allows for parsing the system error log.
†
      catalog = †command.cat†
      setno = 0
      msgno = 42
```

E.6.4 odmget of /etc/objrepos/HACMPevent

```
HACMPevent:
      name = †swap_adapter†
      desc = †Script run to swap IP Addresses between two network adapters.†
      setno = 0
      msgno = 0
      catalog = ††
      cmd = †/usr/sbin/cluster/events/swap_adapter†
      notify = ††
      pre = ††
      post = ††
      recv = ††
      count = 0

HACMPevent:
      name = †swap_adapter_complete†
      desc = †Script run after the swap_adapterscript has successfully completed.†
      setno = 0
      msgno = 0
      catalog = ††
      cmd = †/usr/sbin/cluster/events/swap_adapter_complete†
      notify = ††
      pre = ††
      post = ††
      recv = ††
      count = 0

HACMPevent:
      name = †network_up†
      desc = †Script run after a network has become active.†
      setno = 0
      msgno = 0
      catalog = ††
      cmd = †/usr/sbin/cluster/events/network_up†
      notify = ††
```

```
        pre = ††
        post = ††
        recv = ††
        count = 0

HACMPevent:
        name = †network_down†
        desc = †Script run when a network has failed.†
        setno = 0
        msgno = 0
        catalog = ††
        cmd = †/usr/sbin/cluster/events/network_down†
        notify = ††
        pre = ††
        post = ††
        recv = ††
        count = 0

HACMPevent:
        name = †network_up_complete†
        desc = †Script run after the network_up script has successfully completed.†
        setno = 0
        msgno = 0
        catalog = ††
        cmd = †/usr/sbin/cluster/events/network_up_complete†
        notify = ††
        pre = ††
        post = ††
        recv = ††
        count = 0

HACMPevent:
        name = †network_down_complete†
        desc = †Script run after the network_down script has successfully completed.†
        setno = 0
        msgno = 0
        catalog = ††
        cmd = †/usr/sbin/cluster/events/network_down_complete†
        notify = ††
        pre = ††
        post = ††
        recv = ††
        count = 0

HACMPevent:
        name = †node_up†
        desc = †Script run when a node is attempting to join the cluster.†
        setno = 0
        msgno = 0
        catalog = ††
        cmd = †/usr/sbin/cluster/events/node_up†
        notify = ††
        pre = ††
        post = ††
        recv = ††
```

```
      count = 0

HACMPevent:
      name = †node_down†
      desc = †Script run when a node is attempting to leave the cluster.†
      setno = 0
      msgno = 0
      catalog = ††
      cmd = †/usr/sbin/cluster/events/node_down†
      notify = ††
      pre = ††
      post = ††
      recv = ††
      count = 0

HACMPevent:
      name = †node_up_complete†
      desc = †Script run after the node_up script has successfully completed.†
      setno = 0
      msgno = 0
      catalog = ††
      cmd = †/usr/sbin/cluster/events/node_up_complete†
      notify = ††
      pre = ††
      post = ††
      recv = ††
      count = 0

HACMPevent:
      name = †node_down_complete†
      desc = †Script run after the node_down script has successfully completed.†
      setno = 0
      msgno = 0
      catalog = ††
      cmd = †/usr/sbin/cluster/events/node_down_complete†
      notify = ††
      pre = ††
      post = ††
      recv = ††
      count = 0

HACMPevent:
      name = †join_standby†
      desc = †Script run after a standby adapter has become active.†
      setno = 0
      msgno = 0
      catalog = ††
      cmd = †/usr/sbin/cluster/events/join_standby†
      notify = ††
      pre = ††
      post = ††
      recv = ††
      count = 0

HACMPevent:
```

```
      name = †fail_standby†
      desc = †Script run after a standby adapter has failed.†
      setno = 0
      msgno = 0
      catalog = ††
      cmd = †/usr/sbin/cluster/events/fail_standby†
      notify = ††
      pre = ††
      post = ††
      recv = ††
      count = 0

HACMPevent:
      name = †acquire_service_addr†
      desc = †Script run to configure a service adapter with a service address.†
      setno = 0
      msgno = 0
      catalog = ††
      cmd = †/usr/sbin/cluster/events/acquire_service_addr†
      notify = ††
      pre = ††
      post = ††
      recv = ††
      count = 0

HACMPevent:
      name = †acquire_takeover_addr†
      desc = †Script run to configure a standby adapter with a service address.†
      setno = 0
      msgno = 0
      catalog = ††
      cmd = †/usr/sbin/cluster/events/acquire_takeover_addr†
      notify = ††
      pre = ††
      post = ††
      recv = ††
      count = 0

HACMPevent:
      name = †get_disk_vg_fs†
      desc = †Script run to acquire disks, varyon volume groups, and mount filesystems.†

      setno = 0
      msgno = 0
      catalog = ††
      cmd = †/usr/sbin/cluster/events/get_disk_vg_fs†
      notify = ††
      pre = ††
      post = ††
      recv = ††
      count = 0

HACMPevent:
      name = †node_down_local†
      desc = †Script run when it is the local node which is leaving the cluster.†
```

```
      setno = 0
      msgno = 0
      catalog = tt
      cmd = t/usr/sbin/cluster/events/node_down_localt
      notify = tt
      pre = tt
      post = tt
      recv = tt
      count = 0

HACMPevent:
      name = tnode_down_local_completet
      desc = tScript run after the node_down_local script has successfully completed.t
      setno = 0
      msgno = 0
      catalog = tt
      cmd = t/usr/sbin/cluster/events/node_down_local_completet
      notify = tt
      pre = tt
      post = tt
      recv = tt
      count = 0

HACMPevent:
      name = tnode_down_remotet
      desc = tScript run when it is a remote node which is leaving the cluster.t
      setno = 0
      msgno = 0
      catalog = tt
      cmd = t/usr/HACMP_ANSS/script/CMD_node_down_remotet
      notify = t/usr/HACMP_ANSS/script/event_NOTIFICATIONt
      pre = t/usr/HACMP_ANSS/script/PRE_node_down_remotet
      post = t/usr/HACMP_ANSS/script/POS_node_down_remotet
      recv = tt
      count = 0

HACMPevent:
      name = tnode_down_remote_completet
      desc = tScript run after the node_down_remote script has successfully completed.t
      setno = 0
      msgno = 0
      catalog = tt
      cmd = t/usr/sbin/cluster/events/node_down_remote_completet
      notify = tt
      pre = tt
      post = tt
      recv = tt
      count = 0

HACMPevent:
      name = tnode_up_localt
      desc = tScript run when it is the local node which is joining the cluster.t
      setno = 0
      msgno = 0
      catalog = tt
```

```
      cmd = †/usr/sbin/cluster/events/node_up_local†
      notify = ††
      pre = ††
      post = ††
      recv = ††
      count = 0

HACMPevent:
      name = †node_up_local_complete†
      desc = †Script run after the node_up_local script has successfully completed.†
      setno = 0
      msgno = 0
      catalog = ††
      cmd = †/usr/sbin/cluster/events/node_up_local_complete†
      notify = ††
      pre = ††
      post = ††
      recv = ††
      count = 0

HACMPevent:
      name = †node_up_remote†
      desc = †Script run when it is a remote node which is joining the cluster.†
      setno = 0
      msgno = 0
      catalog = ††
      cmd = †/usr/HACMP_ANSS/script/CMD_node_up_remote†
      notify = †/usr/HACMP_ANSS/script/event_NOTIFICATION†
      pre = †/usr/HACMP_ANSS/script/PRE_node_up_remote†
      post = ††
      recv = ††
      count = 0

HACMPevent:
      name = †node_up_remote_complete†
      desc = †Script run after the node_up_remote script has successfully completed.†
      setno = 0
      msgno = 0
      catalog = ††
      cmd = †/usr/sbin/cluster/events/node_up_remote_complete†
      notify = ††
      pre = ††
      post = ††
      recv = ††
      count = 0

HACMPevent:
      name = †release_service_addr†
      desc = †Script run to configure the boot address on the service adapter.†
      setno = 0
      msgno = 0
      catalog = ††
      cmd = †/usr/sbin/cluster/events/release_service_addr†
      notify = ††
      pre = ††
```

```
      post = ††
      recv = ††
      count = 0

HACMPevent:
      name = †release_takeover_addr†
      desc = †Script run to configure a standby address on a standby adapter.†
      setno = 0
      msgno = 0
      catalog = ††
      cmd = †/usr/sbin/cluster/events/release_takeover_addr†
      notify = ††
      pre = ††
      post = ††
      recv = ††
      count = 0

HACMPevent:
      name = †release_vg_fs†
      desc = †Script run to unmount filesystems and varyoff volume groups.†
      setno = 0
      msgno = 0
      catalog = ††
      cmd = †/usr/sbin/cluster/events/release_vg_fs†
      notify = ††
      pre = ††
      post = ††
      recv = ††
      count = 0

HACMPevent:
      name = †start_server†
      desc = †Script run to start application servers.†
      setno = 0
      msgno = 0
      catalog = ††
      cmd = †/usr/sbin/cluster/events/start_server†
      notify = ††
      pre = ††
      post = ††
      recv = ††
      count = 0

HACMPevent:
      name = †stop_server†
      desc = †Script run to stop application servers.†
      setno = 0
      msgno = 0
      catalog = ††
      cmd = †/usr/sbin/cluster/events/stop_server†
      notify = ††
      pre = ††
      post = ††
      recv = ††
      count = 0
```

```
HACMPevent:
      name = †unstable_too_long†
      desc = †Script run when the Cluster Manger has been unstable for too long.†
      setno = 0
      msgno = 0
      catalog = ††
      cmd = †/usr/sbin/cluster/events/unstable_too_long†
      notify = ††
      pre = ††
      post = ††
      recv = ††
      count = 0

HACMPevent:
      name = †config_too_long†
      desc = †Script run when the Cluster Manger has been in configuration for too long.
†
      setno = 0
      msgno = 0
      catalog = ††
      cmd = †/usr/sbin/cluster/events/config_too_long†
      notify = ††
      pre = ††
      post = ††
      recv = ††
      count = 0

HACMPevent:
      name = †event_error†
      desc = †Script run when a previously executed script has failed to complete succes
sfully.†
      setno = 0
      msgno = 0
      catalog = ††
      cmd = †/usr/sbin/cluster/events/event_error†
      notify = ††
      pre = ††
      post = ††
      recv = ††
      count = 0
```

E.6.5 odmget of /etc/objrepos/HACMPfence

```
HACMPfence:
      pvid = †000009854777a091†
      mask = †0x00000fff†
      nodemap = †goofy:13,mickey:12†

HACMPfence:
      pvid = †000009854777a5c6†
      mask = †0x00000fff†
      nodemap = †goofy:13,mickey:12†
```

E.6.6 odmget of /etc/objrepos/HACMPgroup

```
HACMPgroup:
      group = †mickeyrg†
      type = †cascading†
      nodes = †mickey goofy†

HACMPgroup:
      group = †goofyrg†
      type = †cascading†
      nodes = †goofy mickey†

HACMPgroup:
      group = †concrg†
      type = †concurrent†
      nodes = †mickey goofy†
```

E.6.7 odmget of /etc/objrepos/HACMPnetwork

```
HACMPnetwork:
      name = †etnet1†
      attr = †private†

HACMPnetwork:
      name = †rsnet1†
      attr = †serial†

HACMPnetwork:
      name = †trnet1†
      attr = †public†
```

E.6.8 odmget of /etc/objrepos/HACMPnim

```
HACMPnim:
      name = †ether†
      desc = †Ethernet Protocol†
      addrtype = 0
      path = †/usr/sbin/cluster/nims/nim_ether†
      para = ††
      grace = 30
      hbrate = 500000
      cycle = 12

HACMPnim:
      name = †token†
      desc = †Token Ring Protocol†
      addrtype = 0
      path = †/usr/sbin/cluster/nims/nim_tok†
      para = ††
      grace = 90
      hbrate = 500000
      cycle = 24

HACMPnim:
```

```
        name = †rs232†
        desc = †RS232 Serial Protocol†
        addrtype = 1
        path = †/usr/sbin/cluster/nims/nim_sl†
        para = †† 
        grace = 30
        hbrate = 1500000
        cycle = 6

HACMPnim:
        name = †socc†
        desc = †Serial Optical Protocol†
        addrtype = 0
        path = †/usr/sbin/cluster/nims/nim_socc†
        para = †† 
        grace = 30
        hbrate = 500000
        cycle = 12

HACMPnim:
        name = †fddi†
        desc = †Fiber Data Optical Protocol†
        addrtype = 0
        path = †/usr/sbin/cluster/nims/nim_fddi†
        para = †† 
        grace = 30
        hbrate = 500000
        cycle = 12

HACMPnim:
        name = †IP†
        desc = †Generic IP†
        addrtype = 0
        path = †/usr/sbin/cluster/nims/nim_genip†
        para = †† 
        grace = 30
        hbrate = 500000
        cycle = 12

HACMPnim:
        name = †slip†
        desc = †Serial IP protocol†
        addrtype = 0
        path = †/usr/sbin/cluster/nims/nim_slip†
        para = †† 
        grace = 30
        hbrate = 1000000
        cycle = 12

HACMPnim:
        name = †tmscsi†
        desc = †TMSCSI Serial protocol†
        addrtype = 1
        path = †/usr/sbin/cluster/nims/nim_tms†
        para = †† 
```

```
      grace = 30
      hbrate = 1500000
      cycle = 6

HACMPnim:
      name = †fcs†
      desc = †Fiber Channel Switch†
      addrtype = 0
      path = †/usr/sbin/cluster/nims/nim_fcs†
      para = †t
      grace = 30
      hbrate = 500000
      cycle = 12

HACMPnim:
      name = †hps†
      desc = †High Performance Switch†
      addrtype = 0
      path = †/usr/sbin/cluster/nims/nim_hps†
      para = †t
      grace = 60
      hbrate = 500000
      cycle = 32
```

E.6.9 odmget of /etc/objrepos/HACMPnim.120195

E.6.10 odmget of /etc/objrepos/HACMPnim_pre_U438726

E.6.11 odmget of /etc/objrepos/HACMPnode

```
HACMPnode:
      name = †mickey†
      object = †VERBOSE_LOGGING†
      value = †high†

HACMPnode:
      name = †mickey†
      object = †NAME_SERVER†
      value = †true†

HACMPnode:
      name = †goofy†
      object = †VERBOSE_LOGGING†
      value = †high†

HACMPnode:
      name = †goofy†
      object = †NAME_SERVER†
      value = †true†
```

E.6.12 odmget of /etc/objrepos/HACMPresource

```
HACMPresource:
      group = †mickeyrg†
      name = †SERVICE_LABEL†
      value = †mickey†

HACMPresource:
      group = †mickeyrg†
      name = †FILESYSTEM†
      value = †/test1†

HACMPresource:
      group = †mickeyrg†
      name = †EXPORT_FILESYSTEM†
      value = †/test1†

HACMPresource:
      group = †mickeyrg†
      name = †INACTIVE_TAKEOVER†
      value = †false†

HACMPresource:
      group = †mickeyrg†
      name = †DISK_FENCING†
      value = †false†

HACMPresource:
      group = †mickeyrg†
      name = †SSA_DISK_FENCING†
      value = †false†

HACMPresource:
      group = †goofyrg†
      name = †SERVICE_LABEL†
      value = †goofy†

HACMPresource:
      group = †goofyrg†
      name = †FILESYSTEM†
      value = †/test2†

HACMPresource:
      group = †goofyrg†
      name = †EXPORT_FILESYSTEM†
      value = †/test2†

HACMPresource:
      group = †goofyrg†
      name = †INACTIVE_TAKEOVER†
      value = †false†

HACMPresource:
      group = †goofyrg†
```

```
      name = †DISK_FENCING†
      value = †false†

HACMPresource:
      group = †goofyrg†
      name = †SSA_DISK_FENCING†
      value = †false†

HACMPresource:
      group = †concrg†
      name = †CONCURRENT_VOLUME_GROUP†
      value = †conc1vg†

HACMPresource:
      group = †concrg†
      name = †INACTIVE_TAKEOVER†
      value = †false†

HACMPresource:
      group = †concrg†
      name = †DISK_FENCING†
      value = †false†

HACMPresource:
      group = †concrg†
      name = †SSA_DISK_FENCING†
      value = †false†
```

E.6.13 odmget of /etc/objrepos/HACMPserver

E.6.14 odmget of /etc/objrepos/HACMPsp2

E.6.15 odmget of /etc/objrepos/errnotify

```
errnotify:
      en_pid = 0
      en_name = ††
      en_persistenceflg = 1
      en_label = †CHECKSTOP†
      en_crcid = 0
      en_class = ††
      en_type = ††
      en_alertflg = ††
      en_resource = ††
      en_rtype = ††
      en_rclass = ††
      en_method = †/usr/lib/ras/notifymeth -l $1 -t $9†

errnotify:
      en_pid = 0
      en_name = ††
      en_persistenceflg = 1
      en_label = †CDROM_ERR2†
      en_crcid = 0
```

```
      en_class = ††
      en_type = ††
      en_alertflg = ††
      en_resource = ††
      en_rtype = ††
      en_rclass = ††
      en_method = †/usr/lib/ras/notifymeth -l $1 -r $6 -t $9†

errnotify:
      en_pid = 0
      en_name = ††
      en_persistenceflg = 1
      en_label = †CDROM_ERR4†
      en_crcid = 0
      en_class = ††
      en_type = ††
      en_alertflg = ††
      en_resource = ††
      en_rtype = ††
      en_rclass = ††
      en_method = †/usr/lib/ras/notifymeth -l $1 -r $6 -t $9†

errnotify:
      en_pid = 0
      en_name = ††
      en_persistenceflg = 1
      en_label = †CDROM_ERR6†
      en_crcid = 0
      en_class = ††
      en_type = ††
      en_alertflg = ††
      en_resource = ††
      en_rtype = ††
      en_rclass = ††
      en_method = †/usr/lib/ras/notifymeth -l $1 -r $6 -t $9†

errnotify:
      en_pid = 0
      en_name = ††
      en_persistenceflg = 1
      en_label = †TAPE_ERR3†
      en_crcid = 0
      en_class = ††
      en_type = ††
      en_alertflg = ††
      en_resource = ††
      en_rtype = ††
      en_rclass = ††
      en_method = †/usr/lib/ras/notifymeth -l $1 -r $6 -t $9†

errnotify:
      en_pid = 0
      en_name = ††
      en_persistenceflg = 1
      en_label = †MEMORY†
```

```
      en_crcid = 0
      en_class = tt
      en_type = tt
      en_alertflg = tt
      en_resource = tt
      en_rtype = tt
      en_rclass = tt
      en_method = t/usr/lib/ras/notifymeth -l $1 -t $9t

errnotify:
      en_pid = 0
      en_name = tt
      en_persistenceflg = 1
      en_label = tMEM1t
      en_crcid = 0
      en_class = tt
      en_type = tt
      en_alertflg = tt
      en_resource = tt
      en_rtype = tt
      en_rclass = tt
      en_method = t/usr/lib/ras/notifymeth -l $1 -r $6 -t $9t

errnotify:
      en_pid = 0
      en_name = tt
      en_persistenceflg = 1
      en_label = tMEM2t
      en_crcid = 0
      en_class = tt
      en_type = tt
      en_alertflg = tt
      en_resource = tt
      en_rtype = tt
      en_rclass = tt
      en_method = t/usr/lib/ras/notifymeth -l $1 -r $6 -t $9t

errnotify:
      en_pid = 0
      en_name = tt
      en_persistenceflg = 1
      en_label = tMEM3t
      en_crcid = 0
      en_class = tt
      en_type = tt
      en_alertflg = tt
      en_resource = tt
      en_rtype = tt
      en_rclass = tt
      en_method = t/usr/lib/ras/notifymeth -l $1 -r $6 -t $9t

errnotify:
      en_pid = 0
      en_name = tTAPE_ERR6t
      en_persistenceflg = 1
```

```
      en_label = †TAPE_ERR6†
      en_crcid = 0
      en_class = ††
      en_type = ††
      en_alertflg = ††
      en_resource = ††
      en_rtype = ††
      en_rclass = ††
      en_method = †/usr/lib/ras/notifymeth -l $1 -r $6 -t $9†

errnotify:
      en_pid = 0
      en_name = †sda_err1†
      en_persistenceflg = 1
      en_label = †SDA_ERR1†
      en_crcid = 0
      en_class = †-†
      en_type = †-†
      en_alertflg = †-†
      en_resource = ††
      en_rtype = ††
      en_rclass = ††
      en_method = †/usr/HACMP_ANSS/script/error_SDA $1 $2 $3 $4 $5 $6 $7 $8 $9†

errnotify:
      en_pid = 0
      en_name = †sda_err3†
      en_persistenceflg = 1
      en_label = †SDA_ERR3†
      en_crcid = 0
      en_class = †-†
      en_type = †-†
      en_alertflg = †-†
      en_resource = ††
      en_rtype = ††
      en_rclass = ††
      en_method = †/usr/HACMP_ANSS/script/error_SDA $1 $2 $3 $4 $5 $6 $7 $8 $9†
```

List of Abbreviations

ADSM/6000	Adstar Distributed Storage Manager/6000
AIX	Advanced Interactive Executive
APAR	Authorized Program Analysis Report The description of a problem to be fixed by IBM defect support. This fix is delivered in a PTF (see below).
ARP	Address Resolution Protocol
ASCII	American Standard Code for Information Interchange
AS/400	Application System/400
CDF	Cumulative Distribution Function
CD-ROM	Compact Disk - Read Only Memory
CLM	Cluster Lock Manager
CLVM	Concurrent Logical Volume Manager
CPU	Central Processing Unit
CRM	Concurrent Resource Manager
DE	Differential Ended
DLC	Data Link Control
DMS	Deadman Switch
DNS	Domain Name Service
DSMIT	Distributed System Management Interface Tool
FDDI	Fiber Distributed Data Interface
F/W	Fast and Wide (SCSI)
GB	Gigabyte
GODM	Global Object Data Manager
GUI	Graphical User Interface
HACMP	High Availability Cluster Multi-Processing
HANFS	High Availability Network File System
HCON	Host Connection Program
IBM	International Business Machines Corporation
I/O	Input/Output
IP	Interface Protocol
IPL	Initial Program Load (System Boot)
ITSO	International Technical Support Organization
JFS	Journaled Filesystem
KA	Keepalive Packet
KB	Kilobyte
kb	kilobit
LAN	Local Area Network
LU	Logical Unit (SNA definition)
LUN	Logical Unit (RAID definition)
LVM	Logical Volume Manager
MAC	Medium Access Control
MB	Megabyte
MIB	Management Information Base
MTBF	Mean Time Between Failure
NETBIOS	Network Basic Input/Output System
NFS	Network File System
NIM	Network Interface Module **Note:** This is the definition of NIM in the HACMP context. NIM in the AIX 4.1 context stands for Network Installation Manager.
NIS	Network Information Service
NVRAM	Non-Volatile Random Access Memory
ODM	Object Data Manager
PAD	Packet Assembler/Disassembler
POST	Power On Self Test
PTF	Program Temporary Fix A fix to a problem described in an APAR (see above).
RAID	Redundant Array of Independent (or Inexpensive) Disks
RISC	Reduced Instruction Set Computer
SCSI	Small Computer Systems Interface
SLIP	Serial Line Interface Protocol

SMIT	System Management Interface Tool
SMP	Symmetric Multi-Processor
SMUX	SNMP (see below) Multiplexor
SNA	Systems Network Architecture
SNMP	Simple Network Management Protcol
SOCC	Serial Optical Channel Converter
SPOF	Single Point of Failure
SPX/IPX	Sequenced Package Exchange/Internetwork Packet Exchange
SRC	System Resource Controller
SSA	Serial Storage Architecture
TCP	Transmission Control Protocol
TCP/IP	Transmission Control Protocol/Interface Protocol
UDP	User Datagram Protocol
UPS	Uninterruptible Power Supply
VGDA	Volume Group Descriptor Area
VGSA	Volume Group Status Area
WAN	Wide Area Network

Index

Special Characters

A

B

C

D

E

F

G

U

V

Y

ITSO Technical Bulletin Evaluation RED000

International Technical Support Organization
An HACMP Cookbook
December 1995

Publication No. SG24-4553-00

Your feedback is very important to help us maintain the quality of ITSO Bulletins. **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) If you are an employee of IBM or its subsidiaries:

Do you provide billable services for 20% or more of your time? Yes____ No____

Are you in a Services Organization? Yes____ No____

b) Are you working in the USA? Yes____ No____

c) Was the Bulletin published in time for your needs? Yes____ No____

d) Did this Bulletin meet your needs? Yes____ No____

If no, please explain:

_____

_____

What other topics would you like to see in this Bulletin?

_____

_____

What other Technical Bulletins would you like to see published?

_____

Comments/Suggestions: (THANK YOU FOR YOUR FEEDBACK!)

Name	Address
Company or Organization	
Phone No.	

Fold and Tape **Please do not staple** Fold and Tape

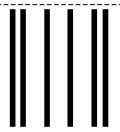

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

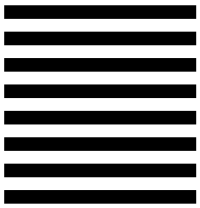

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM International Technical Support Organization
Department JN9B, Building 045
Internal Zip 2834
11400 BURNET ROAD
AUSTIN TX
USA 78758-3493

Fold and Tape **Please do not staple** Fold and Tape

Printed in U.S.A.