

High Availability Cluster Multi-Processing for Linux

HACMP for Linux: Installation and Administration Guide

Version 5.4.1

Second Edition (October 2007)

Before using the information in this book, read the general information in [Notices for HACMP for Linux: Installation and Administration Guide](#).

This edition applies to HACMP for Linux v. 5.4.1, and to all subsequent releases of this product until otherwise indicated in new editions.

© Copyright International Business Machines Corporation, 2006-2009. All rights reserved.

Note to U.S. Government Users Restricted Rights—Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About This Guide	7
Chapter 1: HACMP for Linux Cluster Overview	13
Overview	13
Related Cluster Information	14
Cluster Terminology	15
Sample Configuration with a Diagram	17
Node and Network Failure Scenarios	17
Node Failure	18
Network Failure	18
Network Interface Failure	19
Where You Go from Here	19
Chapter 2: Planning and Installing HACMP for Linux	21
Cluster Hardware	21
Hardware for Cluster Nodes	21
Hardware for Cluster Networks	21
Hardware for Cluster Storage	22
Cluster Software	23
Planning the HACMP Configuration	23
Planning Applications	23
Planning HACMP Nodes	24
Planning for Resource Groups in an HACMP Cluster	25
Resource Group Policies: Overview	26
Planning IP Networks and Network Interfaces	26
Installing HACMP for Linux	27
Software Prerequisites for Installation	28
Contents of the Installation Media	29
Installation Process Overview	30
Installing HACMP for Linux RPMs	30
Installing and Configuring WebSMIT	31
Verifying the Installed Cluster Software	32
Entries Added to System Directories after Installation	33
Addressing Problems during Installation	33
Security Considerations	33
WebSMIT Security	34
Where You Go from Here	34

Chapter 3:	Common Task Summary	35
	Features	35
	Configuration Requirements	36
	Understanding Cluster Network Requirements and Heartbeating	36
	Launching the WebSMIT Interface	37
	Starting WebSMIT	37
	WebSMIT Tasks Overview	37
	Tasks for Configuring a Basic Cluster	38
	Defining a Cluster Name	39
	Adding Nodes and Establishing Communication Paths	40
	Configuring Serial Networks for Heartbeating	41
	Adding IP-Based Networks	41
	Configuring Communication Interfaces/Devices to HACMP	42
	Adding Persistent IP Labels for Cluster Administration Purposes	44
	Verifying the Networks and Network Interfaces Configuration	46
	Configuring Resources to Make Highly Available	48
	Configuring Service IP Labels	48
	Configuring HACMP Application Servers	49
	Configuring Application Servers	49
	Configuring Application Monitors	50
	Including Resources into Resource Groups	50
	Synchronizing the HACMP Cluster Configuration	51
	Displaying the HACMP Cluster Configuration	52
	Viewing the Cluster Status	52
	System Management (C-SPOC) Tasks	54
	Starting HACMP Cluster Services	54
	Stopping HACMP Cluster Services	54
	Managing Resource Groups and Applications	54
	Viewing and Managing Logs	55
	List of Reserved Words	56
	Where You Go from Here	56
Chapter 4:	Other Cluster Configuration Tasks	57
	Identifying Service Adapter Failure for Two-Node Clusters: netmon.cf	57
	Creating the netmon.cf File	57
	Choosing IP Addresses for the netmon.cf File	58
	Resetting Cluster Tunables	59
	Resetting HACMP Tunable Values using the Command Line	59
	Listing Tunable Values	59
	Understanding How HACMP Resets Cluster Tunables	60
	Where You Go from Here	61
Chapter 5:	Monitoring and Troubleshooting a Cluster	63
	Problem Determination Tools	63
	Viewing Cluster Information (clstat) in WebSMIT	63

Useful Commands	64
Logging Messages	64
Collecting Cluster Log Files for Problem Reporting	64
Solving Common Problems with Networks and Applications ...	65
Identifying Causes of Unexpected Network and Network Interface Failures	65
Troubleshooting an Unsuccessful Application Fallover to Another Node	66
Troubleshooting the Serial Connection	66
Appendix A: Command Reference and the clinfo Utility	69
Glossary	71
Index	77

About This Guide

This guide provides information necessary to plan, install, configure, and maintain High Availability Cluster Multi-Processing (HACMP™) on Linux. This guide describes version 5.4.1 of the software.

The following table provides version and manual part numbers for the *HACMP for Linux: Installation and Administration Guide*.

HACMP for Linux Version	Book Name	Book Number
5.4.1	<i>HACMP for Linux: Installation and Administration Guide</i>	SC23-5211-01
5.4 Initial release, 08/2006	<i>HACMP for Linux: Installation and Administration Guide</i>	SC23-5211-00

Who Should Use This Guide

This guide is intended for system administrators and customer engineers responsible for:

- Planning hardware and software resources for an HACMP cluster on the Linux platform.
- Installing and configuring HACMP for Linux.
- Maintaining and troubleshooting HACMP for Linux.

As a prerequisite for installing and configuring the HACMP for Linux package, you should be familiar with:

- Concepts of HACMP (described in this guide).
- Linux operating system: Red Hat™ Enterprise Linux (ES RHEL) v. 4.0 or SUSE™ LINUX Enterprise Server 9.
- RPM Package Manager™ and other Linux utilities.
- Communication protocols, including the TCP/IP subsystem.

Highlighting

This guide uses the following highlighting conventions:

<i>Italic</i>	Identifies new terms or concepts, or indicates emphasis.
Bold	Identifies routines, commands, keywords, files, directories, menu items, and other items whose actual names are predefined by the system.
Monospace	Identifies examples of specific data values, examples of text similar to what you might see displayed, examples of program code similar to what you might write as a programmer, messages from the system, or information that you should actually type.

ISO 9000

ISO 9000 registered quality systems were used in the development and manufacturing of this product.

HACMP Publications

The HACMP software comes with the following publications:

- *HACMP for AIX Release Notes* in **/usr/es/sbin/cluster/release_notes** describe issues relevant to HACMP on the AIX platform: latest hardware and software requirements, last-minute information on installation, product usage, and known issues.
- *HACMP on Linux Release Notes* in **/usr/es/sbin/cluster/release_notes.linux/** describe issues relevant to HACMP on the Linux platform: latest hardware and software requirements, last-minute information on installation, product usage, and known issues.
- *HACMP for AIX: Administration Guide*, SC23-4862
- *HACMP for AIX: Concepts and Facilities Guide*, SC23-4864
- *HACMP for AIX: Installation Guide*, SC23-5209
- *HACMP for AIX: Master Glossary*, SC23-4867
- *HACMP for AIX: Planning Guide*, SC23-4861
- *HACMP for AIX: Programming Client Applications*, SC23-4865
- *HACMP for AIX: Troubleshooting Guide*, SC23-5177
- *HACMP on Linux: Installation and Administration Guide*, SC23-5211
- *HACMP for AIX: Smart Assist Developer's Guide*, SC23-5210
- *IBM International Program License Agreement*.

HACMP/XD Publications

The HACMP Extended Distance (HACMP/XD) software solutions for disaster recovery, added to the base HACMP software, enable a cluster to operate over extended distances at two sites. HACMP/XD publications include the following:

- *HACMP/XD for Geographic LVM (GLVM): Planning and Administration Guide*, SA23-1338
- *HACMP/XD for HAGEO Technology: Concepts and Facilities Guide*, SC23-1922
- *HACMP/XD for HAGEO Technology: Planning and Administration Guide*, SC23-1886
- *HACMP/XD for Metro Mirror: Planning and Administration Guide*, SC23-4863.

HACMP Smart Assist Publications

The HACMP Smart Assist software helps you quickly add an instance of certain applications to your HACMP configuration so that HACMP can manage their availability. The HACMP Smart Assist publications include the following:

- *HACMP Smart Assist for DB2 User's Guide*, SC23-5179
- *HACMP Smart Assist for Oracle User's Guide*, SC23-5178
- *HACMP Smart Assist for WebSphere User's Guide*, SC23-4877
- *HACMP for AIX: Smart Assist Developer's Guide*, SC23-5210
- *HACMP Smart Assist Release Notes*.

IBM AIX Publications

The following publications offer more information about IBM technology related to or used by HACMP:

- *RS/6000 SP High Availability Infrastructure*, SG24-4838
- *IBM AIX v.5.3 Security Guide*, SC23-4907
- *IBM Reliable Scalable Cluster Technology for AIX and Linux: Group Services Programming Guide and Reference*, SA22-7888
- *IBM Reliable Scalable Cluster Technology for AIX and Linux: Administration Guide*, SA22-7889
- *IBM Reliable Scalable Cluster Technology for AIX: Technical Reference*, SA22-7890
- *IBM Reliable Scalable Cluster Technology for AIX: Messages*, GA22-7891.

Accessing Publications

Use the following Internet URLs to access online libraries of documentation:

AIX, IBM eServer Series p[™], and related products:

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

AIX v.5.3 publications:

<http://www.ibm.com/servers/eserver/pseries/library/>

WebSphere Application Server publications:

Search the IBM website to access the WebSphere Application Server Library

DB2 Universal Database Enterprise Server Edition publications:

http://www.ibm.com/cgi-bin/db2www/data/db2/udb/winos2unix/support/v8pubs.d2w/en_main#V8PDF

Tivoli Directory Server publications:

<http://publib.boulder.ibm.com/tividd/td/IBMDirectoryServer5.1.html>

IBM Welcomes Your Comments

You can send any comments via e-mail to hafeedbk@us.ibm.com. Make sure to include the following in your comment or note:

- Title and order number of this book
- Page number or topic related to your comment.

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

Trademarks

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

- AFS
- AIX
- DFS
- **@server**
- eServer Cluster 1600
- Enterprise Storage Server
- HACMP
- IBM
- NetView
- RS/6000
- Scalable POWERParallel Systems
- Series p
- Series x
- Shark
- SP
- WebSphere
- Red Hat Enterprise Linux (RHEL)

- SUSE Linux Enterprise Server
- RPM Package Manager for Linux and other Linux trademarks.

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

Linux is a registered trademark in the United States and other countries and is licensed exclusively through the GNU General Public License.

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

Read the HACMP for Linux Release Notes in the **`/usr/es/sbin/cluster/release.notes_linux`** directory for information that does not appear in the product documentation.

Chapter 1: HACMP for Linux Cluster Overview

High Availability Cluster Multi-Processing (HACMP™) on Linux is the IBM tool for building Linux-based computing platforms that include more than one server and provide high availability of applications and services.

Both HACMP for AIX and HACMP for Linux versions use a common software model and present a common user interface (WebSMIT). This chapter provides an overview of HACMP on Linux and contains the following sections:

- [Overview](#)
- [Cluster Terminology](#)
- [Sample Configuration with a Diagram](#)
- [Node and Network Failure Scenarios](#)
- [Where You Go from Here.](#)

Overview

HACMP for Linux enables your business application and its dependent resources to continue running either at its current hosting server (node) or, in case of a failure at the hosting node, at a backup node, thus providing high availability and recovery for the application.

HACMP detects component failures and automatically transfers your application to another node with little or no interruption to the application's end users.

HACMP for Linux takes advantage of the following software components to reduce application downtime and recovery:

- Linux operating system (RHEL or SUSE ES versions)
- TCP/IP subsystem
- High Availability Cluster Multi-Processing (HACMP™) on Linux cluster management subsystem (the Cluster Manager daemon).

HACMP for Linux provides:

- High Availability for system processes, services and applications that are running under HACMP's control. HACMP ensures continuing service and access to applications during hardware or software outages (or both), planned or unplanned, in an eight-node cluster. Nodes may have access to the data stored on shared disks over an IP-based network (although shared disks cannot be part of the HACMP for Linux cluster and are *not* kept highly available by HACMP).
- Protection and recovery of applications when components fail. HACMP protects your applications against node and network failures, by providing automatic recovery of applications.

If a node fails, HACMP recovers applications on a surviving node. If a network or a network interface card (adapter) fails, HACMP uses an alternate networks, an additional network interface or an IP label alias to recover the communication links and continue providing access to the data.

- WebSMIT, a web-based user interface to configure an HACMP cluster. In WebSMIT, you can configure a basic cluster with the most widely used, default settings, or configure a customized cluster while having the access to customizable tools and functions. WebSMIT lets you view your existing cluster configuration in different ways (node-centric view, or application-centric view) and provides cluster status tools.
- Easy customization of how applications are managed by HACMP. You can configure HACMP to handle applications in the way you want:
 - Applications startup. You select from a set of options for how you want HACMP to start up applications on the node(s).
 - Applications recovery actions that HACMP takes. If a failure occurs with an application's resource that is monitored by HACMP, you select whether you want HACMP to recover applications on another cluster node, or stop the applications.
 - HACMP's follow-up after recovery. You select how you want HACMP to react in cases when you have restored a failed cluster component. For instance, you decide on which node HACMP should restart the application that was previously automatically stopped (or moved to another node) due to a previously detected resource failure.
- Built-in configuration, system maintenance and troubleshooting functions. HACMP has functions to help you with your daily system management tasks, such as cluster administration, automatic cluster monitoring of the application's health, or notification upon component failures.
- Tools for creating similar clusters from an existing "sample" cluster. You can save your existing HACMP cluster configuration in a cluster snapshot file, and later recreate it in an identical cluster in a few steps.

Related Cluster Information

For additional information on IBM on Linux offerings, see:

<http://www.ibm.com/linux/>

<http://www.ibm.com/software/os/linux/software/resource.html>

Cluster Terminology

The list below includes basic terms used in the HACMP environment.

Note: In general, terminology for HACMP is based on industry conventions for high availability. However, the meaning of some of the terms in HACMP may differ from the generic terms.

An application is a service, such as a database, or a collection of system services and their dependent resources, such as a service IP label and application's start and stop scripts, that you want to keep highly available with the use of HACMP.

An application server is a collection of application start and stop scripts that you provide to HACMP by entering the pathnames for the scripts in the WebSMIT user interface. An application server becomes *a resource* associated with an application, you include it in a *resource group* for HACMP to keep it highly available. HACMP ensures that the application can start and stop successfully no matter on which cluster node it is being started.

A cluster node is a physical machine, typically an AIX or a Linux server on which you install HACMP. A cluster node also hosts an application. A cluster node serves as a server for application's clients. HACMP's role is to ensure continuous access to the application, no matter on which node in the cluster the application is currently active.

A home node is a node on which the application is hosted, based on your default configuration for the application's resource group, and under normal conditions.

A takeover node is a backup cluster node to which HACMP may move the application. You can move the application to this node manually, for instance, to free the home node for planned maintenance. Or, HACMP moves the application automatically, due to a cluster component failure.

In HACMP for Linux v.5.4.1, a cluster configuration includes up to eight nodes. Therefore, you can have more than one potential takeover nodes for a particular application. You define the list of nodes on which you want HACMP to host your application using the WebSMIT interface. This list is called a resource group's *nodelist*.

A cluster IP network is used for cluster communications between the nodes and for sending heartbeating information. All IP labels configured on the same HACMP network share the netmask, but may be required to have different subnets.

An IP label is a name of a network interface card (NIC) that you provide to HACMP. Network configuration for HACMP requires planning for several types of IP labels:

- *Base (or boot) IP labels* on each node—the ones through which an initial cluster connectivity is established.
- *Service IP labels* for each application—the ones through which a connection for a highly available application is established.
- Backup IP labels (optional).
- *Persistent IP labels* on each node. These are node-bound IP labels that are useful to have in the cluster for administrative purposes.

Note that to ensure high availability and access to the application, HACMP “recovers” the service IP address associated with the application on another node in the cluster in cases of network interface failures. HACMP uses IP aliases for HACMP networks. For information, see [Planning IP Networks and Network Interfaces](#).

An IP alias is an alias placed on an IP label. It coexists on an interface along with the IP label. Networks that support Gratuitous ARP cache updates enable configuration of IP aliases.

IP Address Takeover (IPAT) is a process whereby a service IP label on one node is taken over by a backup node in the cluster. HACMP uses IPAT to provide high availability of IP service labels that belong to resource groups. These labels provide access to applications. HACMP uses IPAT to recover the IP label on the same node or the backup node. HACMP for Linux by default supports the mode of IPAT known as *IPAT via IP Aliasing*. (The other method of IPAT—*IPAT via IP Replacement* is *not* supported).

IP Address Takeover via IP Aliasing is the default method of IPAT used in HACMP. HACMP uses IPAT via IP Aliasing in cases when it must automatically recover a service IP label on another node. To configure IPAT via IP Aliasing, you configure service IP labels and their aliases to the system. When HACMP performs IPAT during automatic cluster events, it places an IP alias recovered from the “failed” node on top of the service IP address on the takeover node. As a result, access to the application continues to be provided.

Cluster resources can include an application server and a service IP label. All or some of these resources can be associated with an application you plan to keep highly available. You include cluster resources into *resource groups*.

A resource group is a collection of cluster resources.

Resource group startup is an activation of a resource group and its associated resources on a specified cluster node. You choose a *resource group startup policy* from a predefined list in WebSMIT.

Resource group fallover is an action of a resource group, when HACMP moves it from one node to another. In other words, a resource group and its associated application *fall over* to another node. You choose a *resource group fallover policy* from a predefined list in WebSMIT.

Takeover is an automatic action during which HACMP takes over resources from one node and moves them to another node. Takeover occurs when a resource group falls over to another node. A backup node is referred to as a takeover node.

Resource group fallback is an action of a resource group, when HACMP returns it from a takeover node back to the home node. You choose a *resource group fallback policy* from a predefined list in WebSMIT.

Cluster Startup is the starting of HACMP cluster services on the node(s).

Cluster Shutdown is the stopping of HACMP cluster services on the node(s).

Pre- and post-events are customized scripts provided by you (or other system administrators), which you can make known to HACMP and which will be run before or after a particular cluster event. For more information on pre- and post-event scripts, see the chapter on *Planning Cluster Events* in the *HACMP for AIX Planning Guide*.

Sample Configuration with a Diagram

The following configuration includes:

- Node1 and Node2 running Linux
- A serial network
- An IP-based network.

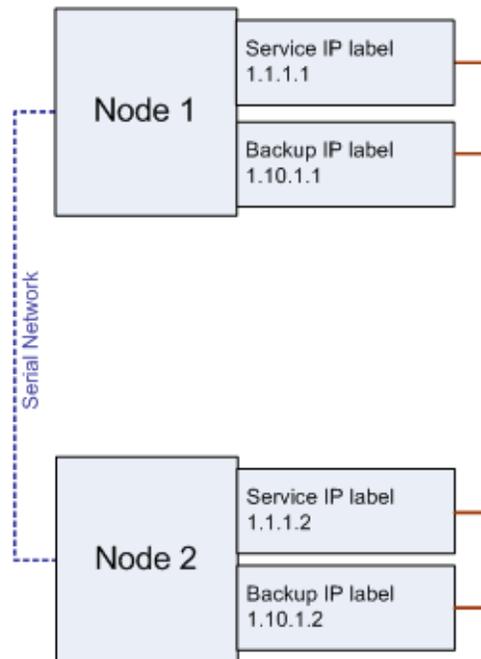


Figure 1. HACMP for Linux Two-Node Cluster Configuration

Node and Network Failure Scenarios

This section describes how HACMP for Linux handles failures and ensures that the application keeps running.

The following scenarios are considered:

- [Node Failure](#)
- [Network Failure](#)
- [Network Interface Failure](#)
- [Preventing Cluster Partitioning.](#)

Node Failure

If the application is configured to normally run on Node1 and Node1 fails, the resource group with the application falls over, or moves, to Node2.

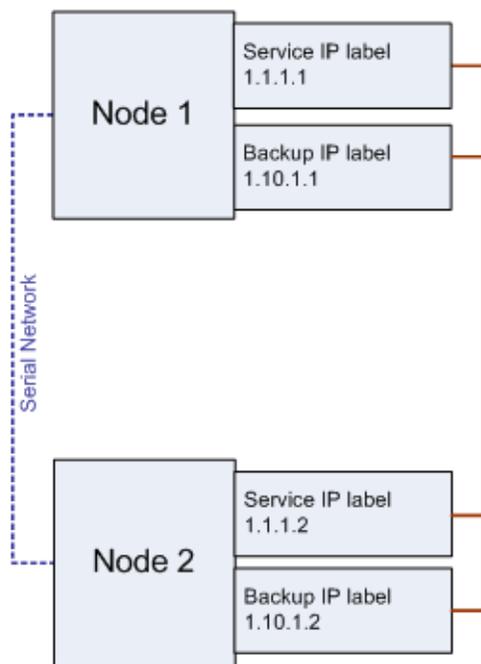


Figure 2. HACMP for Linux Cluster Configuration

At a high level, on Node2, HACMP detects that Node1, the default owner of the resource group, has failed and moves the resource group to Node2. This operation is called a *resource group takeover*. The application is kept highly available and the end users continue to access it.

If Node2 rejoins the cluster, based on the resource group policy HACMP performs the resource group fallback. The resource group moves back to Node1 (for example, if that is the selected fallback policy for the resource group).

Network Failure

A *network failure* occurs when none of the cluster nodes can access each other using any of the network interface cards configured for the HACMP network.

To protect against network failures, we recommend that you have the nodes in the cluster connected by multiple networks. If one network fails, HACMP uses a network that is still available for cluster traffic and for monitoring the status of the nodes (heartbeating).

You can also specify additional actions to process a network failure—for example, re-routing through an alternate network.

How HACMP Handles Network Failures on the Local Node

A local network failure occurs when all interfaces of a specific cluster network on a node fail. For example, if you have nodes A and B, and networks net1 and net2, and all interfaces of network net1 on node A fail, then a **network_down** event runs for net1 with node A as the event node. You can see this in the `/tmp/hacmp.out` file. This is also called a local network failure.

In this case, the Cluster Manager takes selective recovery actions for resource groups containing a service IP label connected to that network. The Cluster Manager attempts to recover only the resource groups affected by the local network failure event.

Network Interface Failure

The HACMP software handles failures of network interfaces on which a service IP label is configured. Types of such failures are:

- Out of two network interfaces configured on the same HACMP node and network, the network interface with a service IP label fails, but an additional “backup” network interface card remains available. In this case, the Cluster Manager removes the service IP label from the failed network interface, and recovers it, via IP aliasing, on the “backup” network interface. Such a network interface failure is transparent to you except for a small delay while the system reconfigures the network interface on the node.
- Out of two network interfaces configured on a node, an additional or a “backup” network interface fails, but the network interface with a service IP label configured on it remains available. In this case, the Cluster Manager detects a (backup) network interface failure, logs the event, and sends a message to the system console. The application continues to be highly available. If you want additional processing, you can customize the processing for this event.
- If the service IP label that is part of a resource group cannot be recovered on a local node, HACMP moves the resource group with the associated IP label to another node, using IP aliasing as the mechanism to recover the associated service IP label.

Preventing Cluster Partitioning

To prevent cluster partitioning, configure a serial network for heartbeating between the nodes, in addition to the IP-based cluster network. If the IP-based cluster network connection between the nodes fails, the heartbeating network prevents data divergence and cluster partitioning.

For information on planning the cluster networks configuration, see [Chapter 2: Planning and Installing HACMP for Linux](#).

For more information, see [Chapter 5: Monitoring and Troubleshooting a Cluster](#).

Where You Go from Here

The remainder of this guide documents how to plan, install, configure and use HACMP for Linux.

1 **HACMP for Linux Cluster Overview** Where You Go from Here

Chapter 2: Planning and Installing HACMP for Linux

This chapter describes how to plan and install HACMP for Linux. It contains the following sections:

- [Cluster Hardware](#)
- [Cluster Software](#)
- [Planning the HACMP Configuration](#)
- [Installing HACMP for Linux](#)
- [Contents of the Installation Media](#)
- [Installation Process Overview](#)
- [Security Considerations](#)
- [Where You Go from Here.](#)

Cluster Hardware

This section lists examples of IBM hardware that you can use for cluster nodes, cluster networks and cluster storage disks. For complete information, see the IBM Portal on Linux website:

<http://www.ibm.com/linux/>

Hardware for Cluster Nodes

HACMP for Linux lets you configure up to eight HACMP nodes. You can use:

- Selected models of IBM System p™ servers

For more information, see: <http://www.ibm.com/systems/p/linux/>

Also, for descriptions of IBM hardware that you can use as HACMP cluster nodes in AIX, see the HACMP for AIX *Planning Guide*.

Hardware for Cluster Networks

HACMP for Linux supports the following interconnection networks for clusters:

- Selected modes of 10/100 Mbps Ethernet
- Selected models of Gigabit Ethernet
- Token Ring.

An Ethernet or a Token Ring network can be used as an HACMP cluster IP-based network.

Hardware for Cluster Storage

HACMP for Linux does not provide high availability for storage resources in your cluster configuration. However you can use NFS or IBM TotalStorage disk subsystems as the storage options in your cluster.

No Automatic NFS and Volume Management

Although you can have disks and file systems configured in the same environment in which your HACMP for Linux cluster is configured, HACMP for Linux does *not* support NFS file systems. You cannot include file systems associated with the application into the resource groups.

This means that the file systems are *not* kept highly available by HACMP for Linux. In particular, during failovers, when applications are moved to other nodes, HACMP for Linux does *not* automatically unmount the associated file systems on one node and mount them on the takeover node. Similarly, HACMP for Linux does *not* automatically perform any volume management or volume group operations for volume groups that a particular application needs to access.

However, if you want to manage storage in the cluster, you can still use NFS or GPFS to control it. To ensure that your NFS file systems work within the cluster, you must manage NFS manually, that is, completely outside of your HACMP for Linux cluster.

For example, for a two-node cluster, you can have an NFS server configured somewhere at your site, and make it to export the file system to your cluster nodes. Both nodes will need to mount the file system at boot time. The file system will be also mounted on another cluster node, the one to which the resource group may potentially fall over in cases of failures. Your application and service IP label will be running on one node. On failover, the application and service IP label will move to the takeover node where the NFS file system has also been mounted since boot time. This way, your application has access to the file system regardless of which node is currently hosting the application. However, the NFS file systems service provided to your application is *not* kept highly available by HACMP.

As an alternative, here is a cluster configuration that lets you have high availability of your NFS file system in the HACMP for Linux cluster. You can configure an NFS server on a separate two-node cluster, with both nodes running HACMP for AIX, specifically, the nodes should run HACMP's NFS component (it is part of HACMP for AIX). You can then export the file system from this highly available NFS server to the nodes of your separate HACMP for Linux cluster.

Cluster Software

The HACMP for Linux cluster software can be described in these two categories:

- Software that you need so that you can install and run the cluster. In particular, HACMP for Linux requires RSCT (IBM Reliable Scalable Cluster Technology) subsystem to be installed on the nodes. For complete information on what software you need to install, see the installation section.
- The application that you plan to make highly available with the use of HACMP. It can be a database or another service. The following link provides an overview of software available for the Linux platform:

<http://www.ibm.com/software/os/linux/software/>

Planning the HACMP Configuration

Plan to have the following components in an HACMP cluster:

- An application
- Up to eight nodes
- Resource groups
- Networks.

Planning Applications

Once you put an application under HACMP's control, HACMP starts it on the node(s) and periodically polls the application's status, if you define application monitors. In cases of component failures, HACMP moves the application to other nodes while the process is invisible to application's end users.

Plan to have the following for your application:

- Customized application start and stop scripts and their locations. The scripts should contain all pre- and post-processing you want HACMP to do so that it starts and stops the applications on the nodes cleanly and according to your requirements. You define these scripts as the application server in WebSMIT.
- Customized scripts you may want to use in HACMP for monitoring the application's successful startup, and for periodically checking the application's running process. You define these scripts to HACMP as application monitors in WebSMIT.
- If you have a complex production environment with tiered applications that require dependencies between their startup, or a staged production environment where some applications should start only if their "supporting" applications are already running, HACMP supports these configurations by letting you configure multiple types of dependencies between resource groups in WebSMIT.

To configure a working cluster that will support such dependent applications, first plan the dependencies for all the services that you want to make highly available. For examples of such planning, see the *HACMP for AIX Planning Guide* and *Administration Guide* (sections on multi-tiered applications and resource group dependencies).

In HACMP 5.4.1, you can use WebSMIT to take an application out of HACMP's control temporarily without disrupting it, and then restart HACMP on the nodes that currently run the application.

Planning HACMP Nodes

HACMP for Linux lets you configure up to eight HACMP nodes.

For each critical application, be mindful of the resources required by the application, including its processing and data storage requirements. For example, when you plan the size of your cluster, include enough nodes to handle the processing requirements of your application after a node fails.

Keep in mind the following considerations when determining the number of cluster nodes and planning the nodes:

- An HACMP cluster can be made up of any combination of supported workstations, LPARs, and other machines. See [Hardware for Cluster Nodes](#). Ensure that all cluster nodes do *not* share components that could be a single point of failure (for example, a power supply). Similarly, do *not* place nodes on a single rack.
- Create small clusters that consist of nodes that perform similar functions or share resources. Smaller, simple clusters are easier to design, implement, and maintain.
- For performance reasons, it may be desirable to use multiple nodes to support the same application. To provide mutual takeover services, the application must be designed in a manner that allows multiple instances of the application to run on the same node.

For example, if an application requires that the dynamic data reside in a directory called */data*, chances are that the application cannot support multiple instances on the same processor. For such an application (running in a non-concurrent environment), try to partition the data so that multiple instances of the application can run—each accessing a unique database.

Furthermore, if the application supports configuration files that enable the administrator to specify that the dynamic data for *instance1* of the application reside in the *data1* directory, *instance2* resides in the *data2* directory, and so on, then multiple instances of the application are probably supported.

- In certain configurations, including additional nodes in the cluster design can increase the level of availability provided by the cluster; it also gives you more flexibility in planning node fallover and reintegration.

The most reliable cluster node configuration is to have at least one standby node.

- Choose cluster nodes that have *enough* I/O slots to support redundant network interface cards and disk adapters.

Ensure you have enough cluster nodes in your cluster. Although this adds to the cost of the cluster, we highly recommend to support redundant hardware, (such as enough I/O slots for network interface cards and disk adapters). This will increase the availability of your application.

- Use nodes with similar processing speed.
- Use nodes with the sufficient CPU cycles and I/O bandwidth to allow the production application to run at peak load. Remember, nodes should have enough capacity to allow HACMP to operate.

To plan for this, benchmark or model your production application, and list the parameters of the heaviest expected loads. Then choose nodes for an HACMP cluster that will *not* exceed 85% busy, when running your production application.

Planning for Resource Groups in an HACMP Cluster

To make your applications highly available in an HACMP cluster, plan and configure resource groups. Resource groups must include resources related to the application, such as its start and stop script (application server) and the service IP label for the application.

Plan the following for resource groups in HACMP for Linux:

- The *nodelist* for the resource groups must contain all or some nodes from the cluster. These are the nodes on which you “allow” HACMP to host your application. The first node in the nodelist is the default node, or the home node for the resource group that contains the application. You define the nodelist in WebSMIT.
- You can use any set of resource group policies for a resource group startup, fallover and fallback. In WebSMIT, HACMP lets you combine only valid sets of these policies and prevents you from configuring non-working scenarios.
- HACMP for Linux supports only non-concurrent resource groups.
- HACMP for Linux does *not* support the fallover policy Fallover using Dynamic Node Priority policy.
- HACMP for Linux does *not* support cluster sites.
- If your applications are dependent on other applications, you may need to plan for dependencies between resource groups. HACMP lets you have node-collocated resource groups, resource groups that always must reside on different nodes, and also child resource groups that do *not* start before their parent resource groups are active (parent/child dependencies). Make a diagram of your dependent applications to better plan dependencies that you want to configure for resource groups, and then define them in WebSMIT.
- HACMP processes the resource groups in parallel by default.
- HACMP for Linux does *not* allow dynamic changes to the cluster resources or resource groups (also known as *dynamic reconfiguration* or *DARE*). This means that you must stop the cluster services, before changing the resource groups or their resources.

For complete planning information, see the guidelines in *Chapter 6: Planning Resource Groups in the HACMP Cluster* in the HACMP for AIX *Planning Guide*.

Resource Group Policies: Overview

HACMP allows you to configure only *valid* combinations of startup, fallover, and fallback behaviors for resource groups. The following table summarizes the basic startup, fallover, and fallback behaviors you can configure for resource groups in HACMP for Linux v. 5.4.1:

Startup Behavior	Fallover Behavior	Fallback Behavior
Online only on home node (first node in the nodelist)	<ul style="list-style-type: none"> Fallover to next priority node in the list 	<ul style="list-style-type: none"> Never fall back or Fall back to higher priority node in the list
Online on first available node	Any of these: <ul style="list-style-type: none"> Fallover to next priority node in the list Bring offline (on error node only) 	<ul style="list-style-type: none"> Never fall back or Fall back to higher priority node in the list

Planning IP Networks and Network Interfaces

Plan to configure the following networks and IP interfaces:

- A heartbeating IP-based network. An HACMP cluster requires at least one network that will be used for the cluster heartbeating traffic.
- A heartbeating serial network, such as RS232.
- An IP-based network that lets you connect from the application's client machine to the nodes. The nodes serve as the application's servers and run HACMP. To configure this network, plan to configure a client machine with a network adapter and a NIC compatible with at least one of the networks configured on the cluster nodes.
- Two HACMP cluster networks. These are TCP/IP-based networks used by HACMP for inter-node communication. HACMP utilities use them to synchronize information between the nodes and propagate cluster changes across the cluster nodes.

For each HACMP cluster network, on each cluster node plan to configure two IP labels that will be available at boot time, will be configured on different subnets, and will be used for IPAT via IP aliasing. See [Planning IP Labels for IPAT via IP Aliasing](#).

- On the cluster node that will serve as a Web server, set up a network connection to access WebSMIT. Typically, you set up WebSMIT to be accessible from the cluster's internal network that is *not* reachable from the Internet. To securely run WebSMIT on a node, you must ensure HTTP(S)/SSL connectivity to that node; it is *not* handled automatically by WebSMIT or HACMP. See [Security Considerations](#).

Planning IP Labels for IPAT via IP Aliasing

IP address takeover via IP aliasing is the default method of taking over the IP address and is supported in HACMP for Linux. IPAT via IP aliasing allows one node to acquire the IP label and the IP address of another node in the cluster, using IP aliases.

To enable that IP Address Takeover via IP aliases can be used in the HACMP for Linux networks configuration, configure NICs for the two HACMP cluster networks that meet the following requirements:

- Plan to configure more than one boot-time IP label on the service network interface card on each cluster node.
- Subnet requirements:
 - Multiple boot-time addresses configured on a node should be defined on different subnets.
 - Service IP addresses must be configured on a different subnet from *all* non-service addresses (such as boot) defined for that network on the cluster node.
- Multiple service labels can coexist as aliases on a given interface.
- The netmask for all IP labels in an HACMP network must be the same.
- Manually add the IP labels described in this section into the `/etc/hosts` file on each node. This must be done before you proceed to configure an HACMP cluster in WebSMIT.

HACMP non-service labels are defined on the nodes as the boot-time addresses, assigned by the operating system after a system boot and before the HACMP software is started. When you start the HACMP software on a node, the node's service IP label is added *as an alias* onto one of the NICs that has a non-service label.

When using IPAT via IP Aliases, the node's NIC must meet the following conditions:

- The NIC has *both* the boot-time and service IP addresses configured, where the service IP label is an alias placed on the interface.
- The boot-time address is never removed from a NIC, simply an alias is added on the NIC in addition to the boot-time address.
- If the node fails, a takeover node acquires the failed node's service address *as an alias* on one of its non-service interfaces on the same HACMP network. During a node fallover event, the service IP label that is moved is placed as an alias on the target node's NIC *in addition* to any other service labels that may already be configured on that NIC.

When using IPAT via IP Aliases, service IP labels are acquired using all available non-service interfaces. If there are multiple interfaces available to host the service IP label, the interface is chosen according to the number of IP labels currently on that interface. If multiple service IP labels are acquired and there are multiple interfaces available, the service IP labels are distributed across all the available interfaces.

Installing HACMP for Linux

HACMP for Linux uses the RPM tool for the installation. For information on RPM, see:

<http://www.rpm.org/>

<http://www.rpm.org/max-rpm/ch-rpm-commands.html>

Once you install HACMP for Linux, proceed to configure WebSMIT for access to the cluster configuration user interface.

Software Prerequisites for Installation

When you install HACMP for Linux, make sure that the following software is installed on the cluster nodes:

- Red Hat™ Enterprise Linux (RHEL) 4 or SUSE™ LINUX Enterprise Server (SLES) 9 (both with latest updates).
Read the readme file for WebSMIT/**usr/es/sbin/cluster/wsm/README** for information on specific Apache V1 and V2 requirements, and for information on specific issues related to RHEL or SUSE Linux distribution.
- RSCT 2.4.5.2. For the latest information about RSCT levels and the latest available APARs for RSCT, check the *HACMP for Linux v. 5.4.1 Release Notes*.
- Apache WebServer V1 and V2 (provided with the Linux distribution).
- ksh93. A compliant version of ksh. Ensure that the **ksh** version you have installed is **ksh93** compliant. The **ksh93** environment is a prerequisite for the RHEL distribution, and HACMP for Linux checks for it prior to the installation.

You can download **ksh93** from the Web. The fileset name is similar to the following:
ksh-20050202-1.ppc.rpm.

Contents of the Installation Media

The HACMP for Linux installation media provides the following **.rpm** files:

<code>hacmp.server-5.4.1.0.ppc.rpm</code>	High Availability Cluster Multi-Processing—server part. <code>hacmp.server</code> provides the server-side functions for HACMP.
---	---

<code>hacmp.client-5.4.1.0.ppc.rpm</code>	High Availability Cluster Multi-Processing—client part. <code>hacmp.client</code> provides the client-side functions for HACMP.
---	---

hacmp.license

<code>hacmp.license-5.4.1.0.ppc.rpm</code>	HACMP for Linux License Package. hacmp.license provides the software License Agreement functions for the HACMP for Linux software.
--	---

hacmp.doc

<code>hacmp.doc.html-5.4.1.0.ppc.rpm</code>	HACMP for Linux HTML documentation—U.S. English
---	---

<code>hacmp.doc.pdf-5.4.1.0.ppc.rpm</code>	HACMP for Linux PDF documentation—U.S. English
--	--

Installation Process Overview

Install the HACMP for Linux software on each cluster node (server). Perform the installation process as the root user.

Installing HACMP for Linux RPMs

Before you install, ensure that you have installed all the prerequisites for the installation. See [Software Prerequisites for Installation](#).

To install HACMP for Linux:

1. Insert the HACMP for Linux CD ROM and install the **hacmp.license.rpm** RPM:

```
rpm -ivh hacmp.license.rpm
```

This RPM provides a utility that lets you accept the License Agreement for HACMP for Linux v.5.4.1, and complete the installation.

Note: You can install the HACMP for Linux documentation without accepting the License Agreement.

2. Run the HACMP installation script **/usr/es/sbin/cluster/install_hacmp**.

This script has two options:

- | | |
|--------------------------------|---|
| -y | Lets you automatically accept the License Agreement. By specifying this flag you agree to the terms and conditions of the License Agreement and will not be prompted. |
| -d <rpm location> | Lets you specify an alternate path to the RPMs for installation, if you are <i>not</i> installing directly from the CD-ROM. |

The **/usr/es/sbin/cluster/install_hacmp** installation script launches the License Agreement Program (LAP) and the License Agreement acceptance dialog appears.

3. Read and accept the License Agreement. The software places a key on your system to identify that you accepted the license agreement.

You can also accept the license without installing the rest of the filesets by running the **/usr/es/sbin/cluster/hacmp_license** script. You can then use the RPM tool to install the remaining RPMs.

The **usr/es/sbin/cluster/install_hacmp** installation script checks for the following prerequisites for the HACMP for Linux software:

- `rsct.basic-x.x.x.x`
- `rsct.core-x.x.x.x`
- `rsct.core.utils-x.x.x.x`
- An appropriate version of `ksh93` (such as `ksh-20050202-1.ppc.rpm`)
- Perl 5 (an RSCT prerequisite. `perl-5.8.3` is installed with RHEL)
- `src-1.3.0.1` (an RSCT prerequisite)

4. Check the required RSCT levels in the *HACMP for Linux v.5.4.1 Release Notes*, or in the section [Software Prerequisites for Installation](#).

You can install HACMP for Linux when prerequisites are already installed, or together with the prerequisites.

The `/usr/es/sbin/cluster/install_hacmp` installation script runs the `rpm` command to install HACMP for Linux RPMs:

```
rpm -ivh hacmp.*
```

5. Verify the installed cluster software. Verify that the RPMs have correct version numbers and other specific information. The RPMs cannot be installed when prerequisites are *not* installed.
6. Configure WebSMIT. See `/usr/es/sbin/cluster/wsm/README` for information, as well as the section [Installing and Configuring WebSMIT](#) in this chapter.
7. Read the *HACMP for Linux v. 5.4.1 Release Notes* `/usr/es/sbin/cluster/release_notes.linux`, for information that does *not* appear in the product documentation.

Note: You can manually install all RPMs without using the `install_hacmp` script.

Installing and Configuring WebSMIT

WebSMIT is a Web-based user interface that provides consolidated access to all functions of HACMP configuration and management, interactive cluster status, and the HACMP documentation.

WebSMIT is:

- Supported on Mozilla-based browsers (Mozilla 1.7.3 for AIX and FireFox 1.5.0.2),
- Supported on Internet Explorer, versions 5.0, 5.5 and 6.0.
- Requires that JavaScript is enabled in your browser.
- Requires network access between the browser and the cluster node that serves as a Web server. To run WebSMIT on a node, you must ensure HTTP(S)/SSL connectivity to that node; it is *not* handled automatically by WebSMIT or HACMP.

To launch the WebSMIT interface:

1. Configure and run a Web server process, such as Apache server, on the cluster node(s) to be administered.
2. See the `/usr/es/sbin/cluster/wsm/README` file for information on basic Web server configuration, the default security mechanisms in place when installing HACMP, and the configuration files available for customization.

You can run WebSMIT on a single node. Note that WebSMIT will be unavailable if a node failure occurs. To provide better availability, you can setup WebSMIT to run on multiple nodes. Since WebSMIT is retrieving and updating information from the HACMP cluster, that information should be available from all nodes in the cluster.

Typically, you set up WebSMIT to be accessible from the cluster's internal network that is *not* reachable from the Internet.

Since the WebSMIT interface runs in a Web browser, you can access it from any platform. For information on WebSMIT security, see [Security Considerations](#).

For more information about installing WebSMIT, see the section *Installing and Configuring WebSMIT* in the *HACMP for AIX Installation Guide*.

Integration of WebSMIT with the Apache Server on Different Linux Distributions

The WebSMIT readme file `/usr/es/sbin/cluster/wsm/README` contains different template files and instructions to enable you to handle variations in packaging, when integrating WebSMIT with the Apache server on different Linux distributions.

Verifying the Installed Cluster Software

After the HACMP for Linux software is installed on all nodes, verify the configuration. Use the verification functions of the RPM utility: your goal is to ensure that the cluster software is the same on all nodes.

Verify that the information returned by the `rpm` command is accurate:

```
rpm -qi hacmp.server
rpm -qi hacmp.client
rpm -qi hacmp.license
rpm -qi hacmp.doc.html
rpm -qi hacmp.doc.pdf
```

Each command should return information about each RPM. In particular, the Name, Version, Vendor, Summary and Description fields should contain appropriate information about each package.

HACMP modifies different system files during the installation process (such as `/etc/inittab`, `/etc/services`, and others). To view the details of the installation process, see the log file `file/var/hacmp/log/hacmp.install.log.<date>`.

Example of the Installation Using RPM

Here is an example of the installation using `rpm`:

```
# rpm -iv hacmp*
Preparing packages for installation...
Cluster services are not active on this node.
hacmp.client-5.4.1.0-06128
Cluster services are not active on this node.
hacmp.server-5.4.1.0-06128
May  8 2006 22:26:18 Starting execution of
/usr/es/sbin/cluster/etc/rc.init
with parameters:

May  8 2006 22:26:18 Completed execution of
/usr/es/sbin/cluster/etc/rc.init
with parameters: .
Exit status = 0
Installation of HACMP for Linux is complete.
```

After installation, use the `rpm` command to view the information about the installed product:

```
ppstest3:~ # rpm -qa | grep hacmp
hacmp.server-5.4.1.0-06128
hacmp.client-5.4.1.0-06128
ppstest3:~ # rpm -qi hacmp.server-5.4.1.0-06128
Name           : hacmp.server           Relocations: (not
relocatable)
Version        : 5.4.1.0           Vendor: IBM Corp.
```

```

Release      : 06128                               Build Date: Mon May  8
21:21:09 2006                                     Build Host:
Install date: Tue May  9 13:03:20 2006           Build Host:
bldlnx18.ppd.pok.ibm.com                         Source RPM:
Group       : System Environment/Base            License: IBM Corp.
hacmp.server-5.4.1.0-06128.nosrc.rpm
Size        : 48627953
Signature   : (none)
Packager    : IBM Corp. <hafeedbk@us.ibm.com>
URL         : http://www.ibm.com/systems/p/ha/
Summary     : High Availability Cluster Multi-Processing - server part
Description :
hacmp.server provides the server side functions for HACMP.

Service information for this package can be found at
  http://techsupport.services.ibm.com/server/cluster

Product ID 5765-G71
Distribution: (none)
  
```

Entries Added to System Directories after Installation

After you install HACMP for Linux, the installation process adds the following lines to the `/usr/es/sbin/cluster/etc/inittab` file:

```
harc:2345:once:/usr/es/sbin/cluster/etc/rc.init >/dev/console 2>&1
```

SRC definitions are added (run `lssrc -s <subsystem>`):

```

Subsystem Group
clcomdES clcomdES
clstrmgrES cluster
topsvcs topsvcs
grpsvcs grpsvcs
  
```

Addressing Problems during Installation

If you experience problems during the installation, refer to the RPM documentation for information on a cleanup process after an unsuccessful installation and other issues.

To view the details of the installation process, see the following log file:

```
/var/hacmp/log/ hacmp.install.log.<date>
```

Security Considerations

HACMP for Linux uses the Cluster Communications utility, `clcomd` to ensure secure communication between the cluster nodes. For complete information on the `clcomd` utility, see the HACMP on AIX *Administration Guide*. `clcomd` in HACMP for Linux does *not* provide the options for message encryption supplied by the CtSec software package from IBM.

Note: WebSMIT relies on its own security utility for all commands that are run with the root privilege. In HACMP for AIX, the WebSMIT security utility uses a set of security APIs specific to AIX, such as updates for the count of failed logins, checks for the expiration of user passwords and checks for who is allowed to login. These commands are *not available* on the HACMP for Linux platform. Use the Pluggable Authentication Module package (PAM) on Linux to provide similar functions.

WebSMIT Security

The network access is required between the browser and the cluster node that serves as a Web server: To run WebSMIT on a node, you must ensure HTTP(S)/SSL connectivity to that node; it is *not* handled automatically by WebSMIT or HACMP.

We recommend using the Internet Protocol SEcurity (IPSEC) technology that provides a secure tunnel into a network, allowing you to connect separate networks through an untrusted connection. Note that IPSEC is *not* included in many Linux distributions.

You may consider using an open source VPN project that will either be used to connect to an existing network that uses the IPSEC standard, or provide a VPN doorway into your network. Check with your Linux distribution provider for availability.

Where You Go from Here

After you install the software on each cluster node, you are ready to configure HACMP for Linux:

- Definitions of cluster terminology and main configuration components are listed in [Chapter 1: HACMP for Linux Cluster Overview](#)
- Overview of basic configuration tasks is included in [Chapter 3: Common Task Summary](#)
- Additional tasks for configuring a cluster in WebSMIT are listed in [Chapter 4: Other Cluster Configuration Tasks](#).

Chapter 3: Common Task Summary

This chapter provides you with a brief overview of the basic tasks needed to configure an HACMP for Linux cluster. It contains the following sections:

- [Features](#)
- [Configuration Requirements](#)
- [Understanding Cluster Network Requirements and Heartbeating](#)
- [Launching the WebSMIT Interface](#)
- [WebSMIT Tasks Overview](#)
- [Tasks for Configuring a Basic Cluster](#)
- [Viewing the Cluster Status](#)
- [System Management \(C-SPOC\) Tasks](#)
- [List of Reserved Words](#)
- [Where You Go from Here.](#)

The HACMP for AIX *Administration Guide* contains a detailed description of WebSMIT components and how to use them. It also describes in detail all the configuration tasks, including the tasks listed in this chapter.

Features

HACMP for Linux has the following features.

It supports:

- Eight nodes in the cluster.
- The heartbeating function that is carried over interfaces configured on IP-based networks and on serial networks. Also, heartbeating over IP aliases is supported.
- Ethernet and TokenRing as IP-based networks.
- RS232 serial network can be configured in the cluster for non-IP heartbeating (highly recommended).
- The HACMP snapshot function. It allows you to take a snapshot of an existing cluster configuration and recreate an identical cluster in a few steps, by applying the snapshot.
- IPAT via IP aliasing is supported by default on IP-based cluster networks.

HACMP for Linux does *not* support:

- Sites in the cluster.
- Concurrent resource groups.
- The AIX LVM (Logical Volume Manager) or NFS file systems as resources in an HACMP cluster. HACMP for Linux does *not* keep file systems or volume groups highly available. However, you can have NFS file systems mounted on all nodes on which your application may be hosted in the cluster.
- Mixed clusters of AIX and Linux, or mixed clusters of different distributions of Linux.
- Dynamic reconfiguration of cluster resources (DARE) is *not* supported. This means that after making changes to the previously configured cluster components for the changes to make effect, you must restart the cluster services.

Configuration Requirements

Before configuring an HACMP for Linux cluster, ensure that the following requirements are met:

- All OS-level network configuration is complete. Note that you may need to reconfigure the networks on your servers so that they comply with HACMP's network requirements. In particular, ensure that the following requirements are met:
 - Install all network cabling to provide the necessary physical connectivity.
 - Configure all network IP interfaces on their boot-time addresses.
 - Ensure that the subnet that you plan to use as a subnet for service IP labels is left unused. That is, none of the IP interfaces are configured to boot to addresses configured on the service subnet. This is required for the IPAT via IP aliasing function of HACMP. You may need to obtain an extra subnet from a networking administrator in your facility.
 - Include *all* IP labels that are defined to HACMP on all cluster nodes in the `/etc/hosts` file on each cluster node. The IP labels must be unique.
- Install and configure the application to be made highly available with the use of HACMP:
 - Write and save the application start and stop scripts on one of the nodes. The application start and stop scripts constitute an application server that you configure in HACMP as a resource.
 - Identify and include the service IP label required by the application server in the `/etc/hosts` file on at least one node (this node can serve as a home node for the resource group in which you will include the application).

Understanding Cluster Network Requirements and Heartbeating

To avoid a single point of failure, the cluster should have more than one network configured. Often the cluster has both IP and non-IP based networks, which allows HACMP to use different heartbeat paths. Use the **Add a Network to the HACMP Cluster** WebSMIT panel to configure HACMP IP and point-to-point networks.

You can use any or all of these methods for heartbeat paths:

- Point-to-point networks
- IP-based networks, including heartbeating using IP aliases.

Launching the WebSMIT Interface

Use WebSMIT to:

- Navigate the running cluster.
- View and customize graphical displays of networks, nodes and resource group dependencies.
- View the status of any connected node (with HACMP cluster services running on the nodes).

Starting WebSMIT

For instructions on integrating WebSMIT with your Apache server, and for launching WebSMIT, see the `/usr/es/sbin/cluster/wsm/README` readme file. It contains sample post-install scripts with variables. Each variable is commented with an explanation of its purpose along with the possible values. You can modify the values of the variables to influence the script behavior.

To start WebSMIT:

1. Using a web browser, navigate to the secure URL of your cluster node, for instance enter the URL similar to the following:

```
https://<server_name>.<domain_name>.com:42267
```

The 42267 is the name of the port for HACMP for Linux. The `<domain_name>` entry is optional, it is only necessary if you are logging in to a server that is *not* part of your local network. The system asks you to log in.

2. Log in to the system and press Continue. WebSMIT starts.

WebSMIT Tasks Overview

The main WebSMIT menu in HACMP for Linux contains the following menu items and tabs:

- **Extended Configuration** to configure your cluster.
- **System Management (C-SPOC)**. *C-SPOC (Cluster Single Point of Control)* is an HACMP function that lets you run HACMP cluster-wide configuration commands from one node in the cluster. In HACMP for Linux, you can use System Management (C-SPOC) to start and stop the cluster services and to move, bring online and bring offline resource groups.
- **Problem Determination Tools**. You can customize cluster verification, view current cluster state, view logs, recover from a cluster event failure, configure error notification methods and perform other troubleshooting tasks.
- **HACMP Documentation**. This is the top-level tab that contains a page with links to all online and printable versions of HACMP documentation, including this guide.

Here is the top-level HACMP for Linux WebSMIT menu:

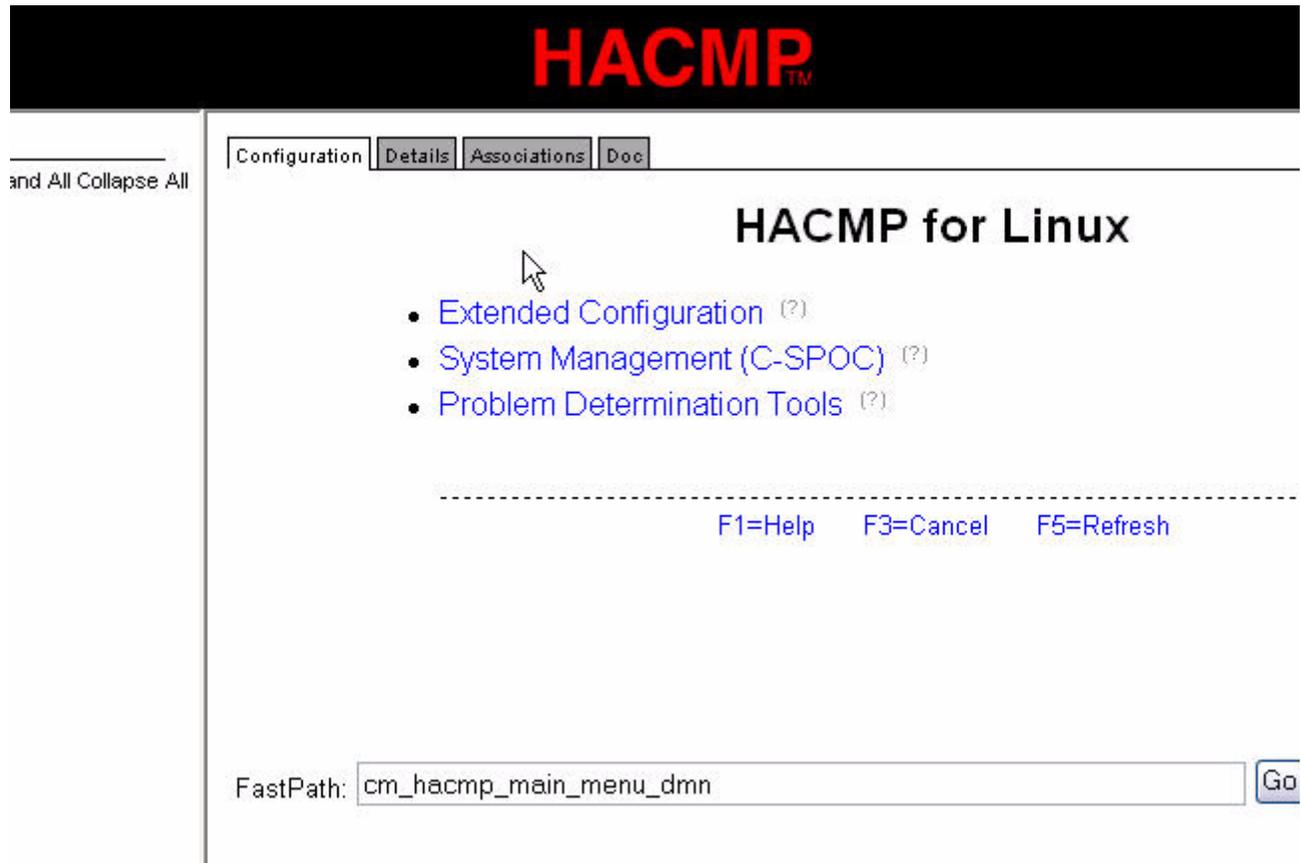


Figure 1. WebSMIT: Top-level Menu

Tasks for Configuring a Basic Cluster

You configure an HACMP for Linux cluster using the Extended Configuration path in WebSMIT.

Note: In general, the sections in this guide provide a high-level overview of user interface options. See the HACMP for AIX *Administration Guide* for detailed procedures, field help, and recommendations for configuring each and every HACMP component.

To configure a basic cluster:

1. On one cluster node, configure a cluster name and add cluster nodes. See:
 - [Defining a Cluster Name](#)
 - [Adding Nodes and Establishing Communication Paths](#)
2. On each cluster node, configure all supporting networks and interfaces: serial networks for heartbeating and IP-based cluster networks for cluster communication.

Also, configure communication devices (that you must have previously defined to the operating system) to HACMP. Configure boot network interfaces (that you must have previously defined to the operating system) to HACMP. Also, configure persistent IP labels for cluster administration purposes. See:

- [Configuring Serial Networks for Heartbeating](#)
 - [Adding IP-Based Networks](#)
 - [Configuring Communication Interfaces/Devices to HACMP](#)
 - [Adding Persistent IP Labels for Cluster Administration Purposes](#)
3. On one cluster node, configure cluster resources that will be associated with the application: service IP labels, application servers and application monitors. See:
 - [Configuring Resources to Make Highly Available](#)
 - [Configuring Service IP Labels](#)
 - [Configuring Application Servers](#)
 - [Configuring Application Monitors](#)
 4. On one cluster node, include resources into resource groups. See [Including Resources into Resource Groups](#).
 5. Synchronize the cluster configuration. See [Synchronizing the HACMP Cluster Configuration](#).
 6. View the HACMP cluster configuration. See [Displaying the HACMP Cluster Configuration](#).
 7. Start the HACMP for Linux cluster services on the cluster nodes. When you do so, HACMP will activate the resource group with the application, and will start monitoring it for high availability. See [Starting HACMP Cluster Services](#).

Defining a Cluster Name

Before starting to configure a cluster:

- Make sure that you added all necessary entries to the `/etc/hosts` file on each machine that will serve as a cluster node. See [Planning IP Networks and Network Interfaces](#).
- Make sure that WebSMIT is installed and can be started on one of the nodes. See [Installing and Configuring WebSMIT](#).
- Log in to WebSMIT. See [Starting WebSMIT](#).

The only step necessary to configure a cluster is to assign the cluster name. When you assign a name to your cluster in WebSMIT, HACMP associates this name with the HACMP-assigned cluster ID.

To assign a cluster name and configure a cluster:

1. Log in to a URL where WebSMIT is installed. The browser window displays the top-level WebSMIT screen.
2. In WebSMIT, select **Extended Configuration > Extended Topology Configuration > Configure an HACMP Cluster > Add/Change/Show an HACMP Cluster** and press Continue.

3. Enter field values as follows:

Cluster Name	Enter an ASCII text string that identifies the cluster. The cluster name can include alphanumeric characters and underscores, but cannot have a leading numeric. Use no more than 32 characters. Do <i>not</i> use reserved names. For a list of reserved names see List of Reserved Words .
---------------------	--

4. Press Continue. If you are changing an existing cluster name, restart HACMP for changes to take effect.

Adding Nodes and Establishing Communication Paths

To add nodes to an HACMP cluster:

1. Log in to a URL where WebSMIT is installed. The browser window displays the top-level WebSMIT screen.
2. In WebSMIT, select **Extended Configuration > Extended Topology Configuration > Configure HACMP Nodes > Add a Node to the HACMP Cluster** and press Continue.
3. Enter field values as follows:

Node name	Enter a unique node name for the node. The name may be up to 32 characters in length. It is <i>not</i> required that the node name be the same as the host name of the node. You can enter one node name at a time, with up to 32 nodes in the cluster.
------------------	---

Communication Path to Node	Enter (or add) one resolvable IP Label (this may be the hostname), IP address, or a Fully Qualified Domain Name for each new node in the cluster, separated by spaces. HACMP uses this path to initiate communication with the node.
-----------------------------------	--

Example 1:

10.11.12.13 <space> NodeC.ibm.com.

Example 2:

NodeA<space>NodeB

(where these are hostnames.)

Or, use the picklist to select the IP labels/addresses that you previously added to **/etc/hosts**.

Once communication paths are established, HACMP adds a new node to the cluster.

4. Repeat this procedure to add up to eight cluster nodes.
5. Proceed to other options in the **Extended Configuration** panel to configure the resources to be made highly available, and assign resources that are to be managed together by HACMP into resource groups.

Configuring Serial Networks for Heartbeating

To configure a serial network that will be used for heartbeating in an HACMP cluster:

1. Log in to a URL where WebSMIT is installed. The browser window displays the top-level WebSMIT screen.
2. In WebSMIT, select **Cluster Configuration and Management > Extended Configuration > Configure HACMP Networks > Add a Network** and press Continue.
3. In the selector screen, select **Predefined Serial Network Types** and the type of network, such as **rs232** and press Continue. The **Add a Serial Network to HACMP Cluster** screen appears.
4. Fill in the fields on the **Add a non IP-based Network** panel as follows:

Network Name	Name the network, using no more than 32 alphanumeric characters and underscores; do <i>not</i> begin the name with a numeric. Do <i>not</i> use reserved names. For a list of reserved names, see List of Reserved Words .
---------------------	--

Network Type	Valid type is RS232 .
---------------------	------------------------------

5. Press Continue to configure this network.
6. Repeat the operation to configure more networks.

For information on testing a serial connection, see [Troubleshooting the Serial Connection](#).

Adding IP-Based Networks

To configure an IP-based cluster network:

1. Log in to a URL where WebSMIT is installed. The browser window displays the top-level WebSMIT screen.
2. In WebSMIT, select **Extended Configuration > Configure HACMP Networks > Add a Network to the HACMP Cluster** and press Continue.
3. In the selector screen, select **Predefined IP-based Network Types**, and then select the type of network, such as **ether** and press Continue. The **Add an IP-Based Network to the HACMP Cluster** window appears.

In the entry fields, values for the interface name *on the local node* are already filled in. HACMP discovers them automatically.

4. Enter the information as follows:

Network Name	If you do <i>not</i> enter a name, HACMP will give the network a default network name made up of the type of network with a number appended (for example, <code>ether1</code>). If you change the name for this network, use no more than 32 alphanumeric characters and underscores.
---------------------	--

Network Type	This field is filled in depending on the type of network you selected.
---------------------	--

Netmask	The netmask, for example, 255.255.255.0.
IP Address Offset for Heartbeating over IP Aliases	<p>(<i>Optional</i>). Leave this field empty if you want HACMP to use the default heartbeating method.</p> <p>Enter a starting address for the heartbeat-over-alias network in this field. For example, 192.168.100.1. You must include the starting address, <i>not</i> just the base network. Do <i>not</i> simply type in 192.168.100 or 192.168.100.0, for example.</p> <p>The network you choose must be on a subnet that is <i>not</i> used by any other network in your physical network setup, and you must have enough available subnets above the one you type in for N networks, where N is the number of interfaces that each node has on the network. Using the example, here, you should have 192.168.100 and 192.168.101 subnets available for an HACMP network that has two interfaces on each node.</p> <p>HACMP uses this address to automatically generate IP addresses for heartbeating, for each boot interface in the configuration. This address range must be unique and must <i>not</i> conflict with any other subnets on the network.</p>

5. Press Continue to configure this network.
6. Repeat the operation to configure more networks.

Configuring Communication Interfaces/Devices to HACMP

To configure network interfaces and devices in HACMP for Linux, ensure that communication interfaces and devices are already configured to the operating system

Now you need to configure them to HACMP. See

- [Configuring Predefined Communication Interfaces to HACMP](#)
- [Configuring Predefined Communication Devices to HACMP](#).

Configuring Predefined Communication Interfaces to HACMP

The **Predefined Communication Interfaces** panel provides fields and picklists that enable you to choose configuration options quickly.

While choosing options, make sure that your choices do *not* conflict with the existing network topology. For example, if your operating system configuration refers to a Token-Ring NIC (Network Interface Card), make sure that HACMP refers to the same type of network interface card (for example, *not* an Ethernet NIC).

Before defining communication interfaces to HACMP, make sure that:

- IP addresses you are defining to HACMP appear in the **ifconfig** output, as in the following example:

```
#ifconfig
eth0      Link encap:Ethernet  HWaddr 00:0D:60:1E:C4:E8
          inet addr:1.1.1.1 Bcast:1.1.1.255  Mask:255.255.255.0
```

- You have entered them in the `/etc/hosts` file on each cluster node.
- If you are adding boot network interfaces to the HACMP configuration, you need to tell HACMP on what NIC the boot interface is configured in the operating system. Use the **ifconfig** command to determine that.

Note that HACMP discovers this information (and presents it to you in a picklist) for the local node, that is, the node on which you are running the WebSMIT session to add interfaces. However, you need to supply this information to HACMP about remote nodes.

To add predefined network interfaces to the HACMP cluster:

1. Log in to a URL where WebSMIT is installed. The browser window displays the top-level WebSMIT screen.
2. In WebSMIT, select **Extended Configuration > Extended Topology Configuration > Configure HACMP Communication Interfaces/Devices > Add Communication Interfaces/Devices** and press Continue.

A window appears that lets you add previously defined network interfaces or devices. Press Continue. WebSMIT displays a selector screen.

3. Select **Communication Interfaces**. The **Select a Network Name** panel appears.
4. Select a network name. The **Add a Communication Interface** panel appears.
5. Fill in the fields as follows:

IP Label/Address	The IP label/address associated with this communication interface that will be configured on the network interface when the node joins the cluster.
Network Type	The type of network media/protocol (for example, ethernet, token ring, etc.) Select the type from the predefined list of network types.
Network Name	A unique name for this logical network.
Node Name	The name of the node on which this network interface physically exists.
Network Interface	<p><i>(Optional)</i> Enter the network interface associated with the communication interface (for example, en0).</p> <p>Note that HACMP discovers the network interface name (and presents it to you in a picklist) for the local node, that is, the node on which you are running the WebSMIT session to add interfaces. However, you need to supply this information about remote nodes to HACMP.</p>

6. Press Continue. You have added the communication interface(s) that were already predefined to the operating system to the HACMP cluster.
7. Verify the configuration by using the `ifconfig` command. For examples, see [Verifying the Networks and Network Interfaces Configuration](#).

Configuring Boot Interfaces

If you are adding boot network interfaces to the HACMP configuration, you need to tell HACMP on what NIC the boot interface is configured in the operating system. Use the `ifconfig` command to determine that.

Note that HACMP discovers this information (and presents it in picklists) on the local node, that is, for the node on which you are running the WebSMIT session to add interfaces. However, you need to supply this information to HACMP about remote nodes.

To add boot network interfaces to the HACMP configuration, use the procedure in [Configuring Predefined Communication Interfaces to HACMP](#).

Configuring Predefined Communication Devices to HACMP

To configure predefined serial devices to the HACMP cluster:

1. Log in to a URL where WebSMIT is installed. The browser window displays the top-level WebSMIT screen.
2. In WebSMIT, select **Extended Configuration > Extended Topology Configuration > Configure HACMP Communication Interfaces/Devices > Add Communication Interfaces/Devices** and press Continue. A panel appears that lets you add previously defined network interfaces or devices.
3. Select the **Communications Devices** and press Continue. WebSMIT displays the **Add a Communications Device** window.
4. Select the non IP-based network to which you want to add the devices.
5. Enter the field values as follows:

Node Name	The node name for the serial device.
Device Name	A device file name. RS232 serial devices must have the device file name <code>/dev/ttyn</code> .
Device Path	For an RS232, for example, <code>/dev/tty0</code>
Network Type	This field is automatically filled in (RS232) when you enter the device name.
Network Name	This field is automatically filled in.

6. Press Continue after filling in all required fields. HACMP now checks the validity of the device configuration. You may receive warnings if a node cannot be reached.
7. Repeat until each node has all appropriate communication devices defined.

Adding Persistent IP Labels for Cluster Administration Purposes

We highly recommend assigning *a persistent node IP label* for a network on a node. A persistent label lets you have a node-bound address on a cluster network that you can use for administrative purposes to access a specific node in the cluster.

For instance, when the HACMP cluster services are running, your previously configured service IP labels may be hosted on any node in the cluster at any given time. Therefore, sending a **ping** command to those labels may *not* result in reaching a specific node. To guarantee access to a particular node, configure a persistent node IP label.

A persistent node IP label is an IP alias that can be assigned to a network for a specified node. A persistent node IP label:

- Always stays on the same node (is node-bound)
- Becomes available at a node's boot time and remains configured on that network when HACMP cluster services are stopped on that node.
- Co-exists with other IP labels present on an interface
- Does *not* require installing an additional physical interface on that node
- *Is not* part of any resource group.

Prerequisites

If you are using persistent node IP labels/addresses, note the following issues:

- Configure only one persistent node IP labels individually on each node per each cluster network.
- On an aliased network (this is the cluster network that is set up to use IPAT via IP Aliasing), a persistent label may be placed on the same subnet as the aliased service label, or it may be configured on an entirely different subnet. However, it must be placed on a different subnet than all boot IP labels on the network.
- You can remove a persistent IP label from the cluster configuration using the **Delete a Persistent Node IP Label/Address** WebSMIT panel. However, after the persistent IP label has been removed from the cluster configuration, it is *not* automatically deleted from the interface on which it was aliased. In order to completely remove the persistent IP label from the node, you should manually remove the alias with the **ifconfig down** command or reboot the cluster node.

To add persistent node IP labels:

1. Log in to a URL where WebSMIT is installed. The browser window displays the top-level WebSMIT screen.
2. In WebSMIT, select **Extended Configuration > Extended Topology Configuration > Configure HACMP Persistent Node IP Labels/Addresses > Add a Persistent Node IP Label** and press Continue.
3. Enter the field values as follows:

Node Name	The name of the node on which the IP label/address will be bound.
Network Name	The name of the network on which the IP label/address will be bound.
Node IP Label/Address	The IP label/address to keep bound to the specified node.

4. Press Continue.

To change or show persistent node IP labels, use the **Change/Show a Persistent Node IP label** WebSMIT menu. To delete them, use the **Delete a Persistent Node IP label** menu.

Verifying the Networks and Network Interfaces Configuration

This section contains several examples of outputs and explains how to verify that you properly configured service IP aliases and netmasks to HACMP:

- [Verifying Service IP Aliases: Example 1](#)
- [Verifying Network Interfaces Configuration: Example 2.](#)

Verifying Service IP Aliases: Example 1

To view the configuration of service IP aliases, use the **ifconfig** command. Here is an example of the output of the **ifconfig** command in HACMP for Linux:

In this example:

- Boot address is 192.9.201.1 configured on the network interface eth2
- Service IP address is 1.1.1.1

In this example, you can see that eth2 is the boot address, and eth2:1 is the service address:

```
eth2      Link encap:Ethernet  HWaddr 00:06:29:B9:1E:7D
          inet addr:192.9.201.1 Bcast:192.9.201.127
Mask:255.255.255.128
          inet6 addr: fe80::206:29ff:feb9:1e7d/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:59153 errors:0 dropped:0 overruns:0 frame:0
          TX packets:18938 errors:0 dropped:0 overruns:0 carrier:0
          collisions:1057 txqueuelen:1000
          RX bytes:12670155 (12.0 Mb)  TX bytes:3363938 (3.2 Mb)
          Interrupt:51 Base address:0xec00

eth2:1    Link encap:Ethernet  HWaddr 00:06:29:B9:1E:7D
          inet addr:1.1.1.1 Bcast:1.255.255.255 Mask:255.255.255.128
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          Interrupt:51 Base address:0xec00
```

To disable a network interface or an alias, run the **ifdown <interface name>** command.

Verifying Network Interfaces Configuration: Example 2

To verify the network interfaces configuration, use the **ifconfig** command, as in the following example:

```
eth0      Link encap:Ethernet  HWaddr 00:06:29:DC:82:CA
          inet addr:192.9.201.3 Bcast:192.9.201.127
Mask:255.255.255.128
          inet6 addr: fe80::206:29ff:fedc:82ca/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:28080 errors:0 dropped:0 overruns:0 frame:0
          TX packets:27980 errors:0 dropped:0 overruns:0 carrier:0
          collisions:3158 txqueuelen:1000
          RX bytes:4624252 (4.4 Mb)  TX bytes:13771958 (13.1 Mb)
          Interrupt:37 Base address:0xec00

eth1      Link encap:Ethernet  HWaddr 00:06:29:B9:1F:71
```

```

    inet addr:192.9.201.133 Bcast:192.9.201.255
Mask:255.255.255.128
    inet6 addr: fe80::206:29ff:feb9:1f71/64 Scope:Link
    UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
    RX packets:5619 errors:1 dropped:0 overruns:0 frame:1
    TX packets:8 errors:0 dropped:0 overruns:0 carrier:0
    collisions:0 txqueuelen:1000
    RX bytes:412380 (402.7 Kb) TX bytes:568 (568.0 b)
    Interrupt:51 Base address:0xec00

lo    Link encap:Local Loopback
    inet addr:127.0.0.1 Mask:255.0.0.0
    inet6 addr: ::1/128 Scope:Host
    UP LOOPBACK RUNNING MTU:16436 Metric:1
    RX packets:213 errors:0 dropped:0 overruns:0 frame:0
    TX packets:213 errors:0 dropped:0 overruns:0 carrier:0
    collisions:0 txqueuelen:0
    RX bytes:15095 (14.7 Kb) TX bytes:15095 (14.7 Kb)

tr0   Link encap:16/4 Mbps Token Ring (New) HWaddr
00:60:94:8A:D3:13
    inet addr:9.57.28.5 Bcast:9.57.28.127 Mask:255.255.255.128
    inet6 addr: fe80::260:94ff:fe8a:d313/64 Scope:Link
    UP BROADCAST RUNNING MULTICAST MTU:4056 Metric:1
    RX packets:9143 errors:0 dropped:0 overruns:0 frame:0
    TX packets:12880 errors:0 dropped:0 overruns:0 carrier:0
    collisions:0 txqueuelen:100
    RX bytes:1959859 (1.8 Mb) TX bytes:2909475 (2.7 Mb)
    Interrupt:38 Base address:0xec00

```

Note that for each NIC there is a line similar to the following:

```
inet addr:192.9.201.3 Bcast:192.9.201.127 Mask:255.255.255.128
```

The address and netmask must match the values you enter in the WebSMIT configuration. After you define the addresses to HACMP, use the **Show Cluster Topology** WebSMIT screen to verify that these values match the values returned by the **ifconfig** command.

After the HACMP cluster services start, the service IP address is added to the NIC and the **ifconfig** output looks similar to the following:

In this configuration:

- The boot address is 192.9.201.1 configured on interface eth2,
- The service IP address is 1.1.1.1

```

eth2   Link encap:Ethernet HWaddr 00:06:29:B9:1E:7D
    inet addr:192.9.201.1 Bcast:192.9.201.127
Mask:255.255.255.128
    inet6 addr: fe80::206:29ff:feb9:1e7d/64 Scope:Link
    UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
    RX packets:59153 errors:0 dropped:0 overruns:0 frame:0
    TX packets:18938 errors:0 dropped:0 overruns:0 carrier:0
    collisions:1057 txqueuelen:1000
    RX bytes:12670155 (12.0 Mb) TX bytes:3363938 (3.2 Mb)
    Interrupt:51 Base address:0xec00

eth2:1 Link encap:Ethernet HWaddr 00:06:29:B9:1E:7D
    inet addr:1.1.1.1 Bcast:1.255.255.255 Mask:255.255.255.128
    UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
    Interrupt:51 Base address:0xec00

```

Configuring Resources to Make Highly Available

In the WebSMIT menu **Configure Resources to Make Highly Available**, you configure the system's resources to be known in HACMP. These are the resources that are associated with your application.

You configure resources that are to be shared among the nodes in the cluster, and kept highly available with HACMP, so that if one component fails, another component automatically takes its place.

Configure these resources (in any order):

- **Service IP label(s)**. This is the service IP label that an application's clients use to connect to the application. HACMP ensures the availability of this IP label by recovering this IP label on other nodes as an alias. The IP label recovery is done with IPAT via IP Aliasing function. See [Configuring Service IP Labels](#).
- **Application servers** (collections of application start and stop scripts) for the application you are making highly available. You provide scripts for starting and stopping your applications to HACMP. Have these scripts ready and specify their pathnames to HACMP. See [Configuring Application Servers](#).
- **Application monitors** for your application. See [Configuring Application Monitors](#).

Configuring Service IP Labels

To add service IP labels/addresses as resources to the resource group in your cluster:

1. Log in to a URL where WebSMIT is installed. The browser window displays the top-level WebSMIT screen.
2. In WebSMIT, select **Extended Configuration > Extended Resource Configuration > HACMP Extended Resources Configuration > Configure HACMP Service IP Labels/Addresses > Add a Service IP Label/Address** and press Continue.
3. Fill in field values as follows:

IP Label/Address	Enter, or select from the picklist the IP label/address to be kept highly available.
Network Name	Enter the symbolic name of the HACMP network on which this Service IP label/address will be configured.

4. Press Continue after filling in all required fields.
5. Repeat the previous steps until you have configured all service IP labels/addresses for each network, as needed.

Configuring HACMP Application Servers

An *application server* is a cluster component that is included in the resource group as a cluster resource, and that is used to control an application that must be kept highly available. An application server consists of application start and stop scripts. Configuring an application server does the following:

- Associates a meaningful name with the server application. For example, you could give the tax software a name such as *taxes*. You then use this name to refer to the application server when you define it as a resource. When you set up the resource group, you add an application server as a resource.
- Points the cluster event scripts to the scripts that they call to start and stop the server application.
- Allows you to then configure application monitoring for that application server.

Note that this section does *not* discuss how to write the start and stop scripts. See the vendor documentation for specific product information on starting and stopping a particular application.

Configuring Application Servers

To configure an application server on any cluster node:

1. Log in to a URL where WebSMIT is installed. The browser window displays the top-level WebSMIT screen.
2. In WebSMIT, select **Extended Configuration > Extended Resource Configuration > HACMP Extended Resources Configuration > Configure HACMP Applications > Configure HACMP Application Servers > Add an Application Server** and press Continue.

WebSMIT displays the **Add an Application Server** panel.

3. Enter field values as follows:

Server Name	Enter an ASCII text string that identifies the server. You will use this name to refer to the application server when you define resources during node configuration. The server name can include alphabetic and numeric characters and underscores. Use no more than 64 characters.
Start Script	Enter the pathname of the script (followed by arguments) called by the cluster event scripts to start the application server. (Maximum 256 characters.) This script must be in the same location on each cluster node that might start the server. The contents of the script, however, may differ.
Stop Script	Enter the pathname of the script called by the cluster event scripts to stop the server. (Maximum 256 characters.) This script must be in the same location on each cluster node that may start the server. The contents of the script, however, may differ.

4. Press Continue to add this information to the HACMP Configuration Database on the local node.
5. Add the application start, stop and notification scripts to every node in the cluster.

Verifying Application Servers

Make sure that the application start, stop and notification scripts exist and are executable on every node in the cluster. Use the **cllsserv** command.

For example:

```
ppstest2:~ # /usr/es/sbin/cluster/utilities/cllsserv
app_test2_primary /usr/local/app_start /usr/local/app_stop
ppstest2:~ # ls -l /usr/local/app_start
-rwxr--r-- 1 root root 169 May 10 22:54 /usr/local/app_start
```

Configuring Application Monitors

Once you configured application servers, HACMP for Linux lets you have application monitors that will check the health of the running application process, or check for the successful start of the application.

To configure application monitors:

1. Log in to a URL where WebSMIT is installed. The browser window displays the top-level WebSMIT screen.
2. In WebSMIT, select **Extended Configuration > Extended Resource Configuration > HACMP Extended Resources Configuration > Configure HACMP Applications > Configure HACMP Application Monitoring** and press Continue. A selector screen appears for **Configure Process Application Monitoring** and **Configure Custom Application Monitoring**.
3. Select the type of monitoring you want and press Continue.
4. Select the application server to which you want to add a monitor.
5. Fill in the field values and press Continue.

For additional reference information on application monitoring, its types, modes, and other information, see the HACMP for AIX *Administration Guide*.

Including Resources into Resource Groups

Once you configure resources to HACMP, you include them in resource groups so that HACMP can manage them as a single set. For example, if an application depends on the service IP label, you can add it to a single resource group.

HACMP manages the resources in a resource group by bringing the resource groups online and offline on their home node(s), or moving them to other nodes, if necessary for recovery.

Note: For detailed instructions on resource groups, see the HACMP for AIX *Administration Guide*. This guide contains descriptions of procedures in HACMP SMIT, and the options are identical to those used in WebSMIT in HACMP for Linux.

Resource Group Management: Overview

In the **Extended Configuration > Extended Resource Configuration > HACMP Extended Resource Group Configuration** WebSMIT screen, you can:

- Add a resource group.
- Change/Show a resource group. The system displays all previously defined resource groups. After selecting a particular resource group, you can view and change the group name, node relationship, and participating nodes (nodelist). You can also change the group's startup, fallover and fallback policies.
- Remove a resource group.
- Change/Show resources for a resource group. Add resources, such as a service IP label for the application, or an application server, to a resource group. HACMP always activates and brings offline these resources on a particular node as a single set. If you want HACMP to activate one set of resources on one node and another set of resources on another node, create separate resource groups for each set.
- Show all resources by node for a resource group.

HACMP for Linux does *not* allow to change resources dynamically, that is, when HACMP cluster services are running on the nodes. To change the previously added resources, stop the cluster services.

Adding Resources to a Resource Group

To include resources into a resource group:

1. Log in to a URL where WebSMIT is installed. The browser window displays the top-level WebSMIT screen.
2. In WebSMIT, select **Extended Configuration > Extended Resource Configuration > HACMP Extended Resources Configuration > Change/Show All Resources and Attributes for a Resource Group** and press Continue.
3. Fill in the field values and press Continue. HACMP adds the resources.

For additional information on adding or changing resources in resource groups, and for information on other resource group management tasks, see the *Administration Guide*.

Synchronizing the HACMP Cluster Configuration

We recommend that you do all the configuration from one node and synchronize the cluster to propagate this information to other nodes.

Use this WebSMIT option to commit and distribute your changes automatically to all of the specified nodes.

To synchronize an HACMP cluster configuration:

1. Log in to a URL where WebSMIT is installed. The browser window displays the top-level WebSMIT screen.
2. In WebSMIT, select **Extended Configuration > Extended Verification and Synchronization** and press Continue.

If you configured the cluster correctly, HACMP synchronizes the configuration. HACMP issues errors if the configuration is *not* valid.

Displaying the HACMP Cluster Configuration

You can ask HACMP to show you the status of different configured components. The WebSMIT options for displaying different cluster entities are grouped together with the options for adding them to the cluster.

Here are some examples of the options you have:

- Show HACMP Topology by node, by network name, or by communication interface
- Change/Show Persistent IP Labels
- Show Cluster Applications and change/show application monitors per application
- Change/Show Service IP Labels
- Show all Resources by Node or Resource Groups
- View cluster logs (In WebSMIT, it is under **System Management > Log Viewing and Management**)
- Show Cluster Services (whether running or *not*).

Viewing the Cluster Status

HACMP has a cluster status utility, the `/usr/es/sbin/cluster/clstat`. It reports the status of key cluster components—the cluster itself, the nodes in the cluster, the network interfaces connected to the nodes, and the resource groups on each node.

`clstat` is available in WebSMIT at the left side of the top-level menu. It displays an expandable list of cluster components along with their status. The cluster status display window shows information and status (**up** or **down**, **online**, **offline** or **error**) on cluster nodes, networks, interfaces, application servers and resource groups. For resource groups, it also shows the node on which the group is currently hosted.

Here is an example of the **clstat** output in WebSMIT. This is the left-hand side panel of the window:

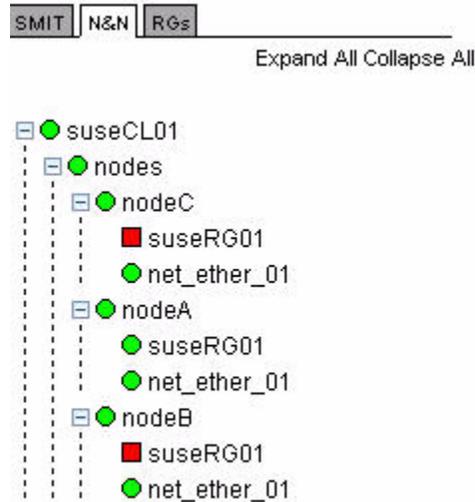


Figure 2. clstat Output

Here is an example of the ASCII-based output from the **clstat** command, used on a Linux cluster with nodes named ppstest1 and ppstest2:

```
ppstest2:~ # /usr/es/sbin/cluster/clstat

clstat - HACMP Cluster Status Monitor
-----
Cluster: test1234          (1148058900)
Wed May 17 16:45:41 2006
State: UP                  Nodes: 4
SubState: STABLE

Node: ppstest1            State: UP
  Interface: tr0 (6)      Address: 9.57.28.3
                          State: UP
  Resource Group: rg1     State: On line

Node: ppstest2            State: UP
  Interface: tr0 (6)      Address: 9.57.28.4
                          State: UP
  Resource Group: rg2     State: On line

Node: ppstest3            State: UP
  Interface: tr0 (6)      Address: 9.57.28.5
                          State: UP

Node: ppstest4            State: UP
  Interface: tr0 (6)      Address: 9.57.28.6
                          State: UP
  Resource Group: rg3     State: On line
  Resource Group: rg4     State: On line
```

System Management (C-SPOC) Tasks

Use the System Management (C-SPOC) panel in WebSMIT to configure *from one node* the resources that are shared among nodes. System Management utility of HACMP lets you administer many aspects of the cluster and its components from one Cluster Single Point of Control (C-SPOC). By automating repetitive tasks, C-SPOC eliminates a potential source of errors, and speeds up the cluster maintenance process.

In WebSMIT, you access C-SPOC using the **System Management (C-SPOC)** menu.

In this panel, you can do the following tasks from one node:

- Manage HACMP services, or start and stop cluster services: Cluster Manager (**clstrmgr**) and Cluster Information (**clinfo**).
- HACMP Communication Interface Management. Manage the communication interfaces of existing cluster nodes using C-SPOC.
- HACMP Resource Group and Application Management Provides menus to manage cluster resource groups and analyze cluster applications.
- HACMP Log Viewing and Management. Manage, view, and collect HACMP log files and event summaries.

Starting HACMP Cluster Services

To start HACMP cluster services:

1. Log in to a URL where WebSMIT is installed. The browser window displays the top-level WebSMIT screen.
2. In WebSMIT, select **System Management (C-SPOC) > Manage HACMP Services > Start HACMP Services** and press Continue.

For detailed instructions, see the HACMP on AIX *Administration Guide*.

Stopping HACMP Cluster Services

To stop HACMP cluster services:

1. Log in to a URL where WebSMIT is installed. The browser window displays the top-level WebSMIT screen.
2. In WebSMIT, select **System Management (C-SPOC) > Manage HACMP Services > Start HACMP Services** and press Continue.

For detailed instructions, see the HACMP on AIX *Administration Guide*.

Managing Resource Groups and Applications

To manage resource groups and applications:

1. Log in to a URL where WebSMIT is installed. The browser window displays the top-level WebSMIT screen.
2. In WebSMIT, select **System Management (C-SPOC) > HACMP Resource Group and Application Management** and press Continue.

For detailed instructions, see the HACMP on AIX *Administration Guide*.

Viewing and Managing Logs

To view and manage logs:

1. Log in to a URL where WebSMIT is installed. The browser window displays the top-level WebSMIT screen.
2. In WebSMIT, select **System Management (C-SPOC) > HACMP Log Viewing and Management** and press Continue.

For detailed instructions, see the HACMP on AIX *Administration Guide*.

List of Reserved Words

Do *not* use the following words as names in a cluster. However, you may use these words when combined with numerals or another word (for example, `my_network` or `rs232_02`).

adapter	false	nim	socc
alias	FBHPN	node	subnet
all	fcs	nodename	tm SCSI
ALL	fddi	OAAN	tmssa
ANY	FNPN	OFAN	token
atm	fscsi	OHN	true
BO	FUDNP	OTHER	tty
cluster	grep	OUDP	volume
command	group	private	vpath
CROSS_SITE_RG_MOVE	hps	public	vscsi
custom	ib	resource	XD_data
daemon	ip	RESTORE	XD_ip
disk	IP	root	XD_rs232
diskhb	name	rs232	
ether	network	serial	
event	NFB	slip	

Where You Go from Here

See the next chapter for information on additional cluster configuration tasks.

Chapter 4: Other Cluster Configuration Tasks

This chapter provides you with a brief overview of additional tasks you may need to do in the HACMP for Linux cluster.

Note: The HACMP for AIX *Administration Guide* contains a detailed description of WebSMIT components and how to use them. It also describes in detail all the configuration tasks.

This chapter describes:

- [Identifying Service Adapter Failure for Two-Node Clusters: netmon.cf](#)
- [Resetting Cluster Tunables](#)
- [Where You Go from Here.](#)

Identifying Service Adapter Failure for Two-Node Clusters: netmon.cf

In cluster configurations where there are networks that under certain conditions can become single adapter networks, it can be difficult for HACMP to accurately determine adapter failure. This is because RSCT Topology Services *cannot* force packet traffic over the single adapter to confirm its proper operation. (This shortcoming is less of an exposure if the network adapter is under heavy use. In this case, the inbound packet count continues to increase over the service adapter without stimulation from RSCT Topology Services).

An enhancement to **netmon**, the network monitor portion of RSCT Topology Services, allows for a more accurate determination of a service adapter failure. This function can be used in configurations that require a single service adapter per network.

You can create **netmon.cf** configuration file that specifies additional network addresses to which ICMP ECHO requests can be sent.

This file must exist at cluster startup—RSCT Topology Services scans the **netmon.cf** configuration file during initialization. When **netmon** needs to stimulate the network to ensure adapter function, it sends ICMP ECHO requests to each IP address. After sending the request to every address, **netmon** checks the inbound packet count before determining whether an adapter has failed.

Creating the netmon.cf File

The **netmon.cf** file must be placed in the `/usr/sbin/cluster` directory on all cluster nodes.

Requirements for creating the file:

- The **netmon.cf** file consists of one IP address or IP label per cable.

- Include each IP address and its corresponding label for the **netmon.cf** file in the `/etc/hosts` file.
- When selecting IP addresses (or hostnames) for the **netmon.cf** file, keep in mind ALL possible IP addresses that an interface might hold at any given time (boot IP addresses, service IP addresses used by HACMP, and other interfaces), and ensure that for each interface on the node, the **netmon.cf** file contains one or more targets that can be reached by those addresses.

For example, an adapter on a node that serves as a boot adapter when it is *not* holding a service address should be able to ping some targets in the **netmon.cf** file using that boot address.

Note: Ensure that the names in the **netmon.cf** file are included in the `/etc/hosts` file. If this is *not* done, then when the NIM process (that is part of the RSCT Topology Services subsystem) attempts to determine the state of the local adapters, it may try to run the hostname resolution. The hostname resolution may in turn result in the process being blocked in cases when the network used for name resolution is unreachable. The blockage may result in longer adapter failure detection times, which will slow failover operations.

- A maximum of 30 IP addresses/labels can be defined in **netmon.cf**.

The following example shows a `/usr/sbin/cluster/netmon.cf` configuration file:

```
180.146.181.119
steamer
chowder
180.146.181.121
mussel
```

Choosing IP Addresses for the `netmon.cf` File

Guidelines for choosing the IP addresses to include in the **netmon.cf** file depend on whether the local interface is a service or a non-service interface:

If the local interface is a service interface, it will verify that it is operational via point-to-point communication with the following interfaces:

- Existing remote service interface(s) on the same logical network
- One of the existing local non-service interfaces on the same logical network
- IP addresses/hostnames in the **netmon.cf** file.

If the local interface is a non-service interface, it verifies that it is operational via point-to-point communication with the following interfaces:

- Existing remote non-service(s) on the same logical network
- One of the existing local interfaces on the same logical network
- IP addresses/hostnames in the **netmon.cf** file.

The **netmon.cf** file should contain remote IP labels/addresses that are *not* in the cluster configuration that can be accessed from HACMP interfaces. Routers can also be used.

Resetting Cluster Tunables

You can change the settings for a list of tunable values that were altered during cluster maintenance and reset them to their default settings, or installation-time cluster settings. The *installation-time* cluster settings are equal to the values that appear in the cluster after installing HACMP from scratch.

Note: Resetting the tunable values *does not* change any other aspects of the configuration, while installing HACMP removes all user-configured configuration information including nodes, networks, and resources.

To reset the cluster tunable values:

1. Stop the cluster services.
2. Log in to a URL where WebSMIT is installed. The browser window displays the top-level WebSMIT screen.
3. In WebSMIT, select **Extended Configuration > Extended Topology Configuration > Configure an HACMP Cluster > Reset Cluster Tunables** and press Continue.

Use this option to reset all the tunables (customizations) made to the cluster. For a list of the tunable values that will change, see the section [Listing Tunable Values](#). Using this option returns all tunable values to their default values but does *not* change the cluster configuration. HACMP takes a snapshot file before resetting. You can choose to have HACMP synchronize the cluster when this operation is complete.

4. Select the options as follows and press Continue:

Synchronize Cluster Configuration If you set this option to **yes**, HACMP synchronizes the cluster after resetting the cluster tunables.

5. HACMP asks: "Are you sure?"
6. Press Continue.

HACMP resets all the tunable values to their original settings and removes those that should be removed (such as the nodes' knowledge about customized pre- and post-event scripts).

Resetting HACMP Tunable Values using the Command Line

We recommend that you use the SMIT interface to reset the cluster tunable values. The **clsnapshot -t** command also resets the cluster tunables. This command is intended for use by IBM support. See the man page for more information.

Listing Tunable Values

You can change and reset the following list of tunable values:

- User-supplied information.
 - Network module tuning parameters, such as, failure detection rate, grace period and heartbeat rate. HACMP resets these parameters to their installation-time default values.

- Cluster event customizations, such as, all changes to cluster events. Note that resetting changes to cluster events does *not* remove any files or scripts that the customization use; it only removes the knowledge HACMP has of pre- and post-event scripts.
- Cluster event rule changes made to the event rules database are reset to the installation-time default values.
- HACMP command customizations made to the default set of HACMP commands are reset to the installation-time defaults.
- Automatically generated and discovered information.
Generally users cannot see this information. HACMP rediscovers or regenerates this information when the cluster services are restarted or during the next cluster synchronization.

HACMP resets the following:

- Local node names stored in the cluster definition database
- Netmasks for all cluster networks
- Netmasks, interface names and aliases for disk heartbeating (if configured) for all cluster interfaces
- SP switch information generated during the latest **node_up** event (this information is regenerated at the next **node_up** event)
- Instance numbers and default log sizes for the RSCT subsystem.

Understanding How HACMP Resets Cluster Tunables

HACMP resets tunable values to their default values under the following conditions:

- Before resetting HACMP tunable values, HACMP takes a cluster snapshot. After the values have been reset to defaults, if you want to go back to your customized cluster settings, you can restore them with the cluster snapshot. HACMP saves snapshots of the last ten configurations in the default cluster snapshot directory, **/usr/es/sbin/cluster/snapshots**, with the name **active.x.odm**, where x is a digit between 0 and 9, with 0 being the most recent.
- Stop cluster services on all nodes before resetting tunable values. HACMP prevents you from resetting tunable values in a running cluster.

In some cases, HACMP cannot differentiate between user-configured information and discovered information, and *does not* reset such values. For example, you may enter a service label and HACMP automatically discovers the IP address that corresponds to that label. In this case, HACMP does *not* reset the service label or the IP address. The cluster verification utility detects if these values do *not* match.

The **clsnapshot.log** file in the snapshot directory contains log messages for this utility. If any of the following scenarios are run, then HACMP cannot revert to the previous configuration:

- **cl_convert** is run automatically
- **cl_convert** is run manually

Where You Go from Here

The following chapters describe:

- Logging and troubleshooting utilities
- Interaction of HACMP with other client applications (**clinfo** utility)
- Command reference
- Glossary of HACMP terms.

4 **Other Cluster Configuration Tasks**

Where You Go from Here

Chapter 5: Monitoring and Troubleshooting a Cluster

This chapter presents general information for monitoring and troubleshooting an HACMP for Linux configuration.

This chapter contains the following sections:

- [Problem Determination Tools](#)
- [Viewing Cluster Information \(clstat\) in WebSMIT](#)
- [Useful Commands](#)
- [Logging Messages](#)
- [Solving Common Problems with Networks and Applications.](#)

Problem Determination Tools

WebSMIT Problem Determination Tools menu has a set of tools for troubleshooting and recovering from problems that may arise in a cluster environment.

The **Problem Determination Tools** panel in WebSMIT includes:

- **View Current State.** WebSMIT displays cluster information using a slightly different layout and organization. Cluster components are displayed along their status. Expanding the item reveals additional information about it, including network, interfaces and active resource groups.
- **HACMP Log Viewing and Management.** Contains utilities that display or manage logs maintained by HACMP. These include the log file named **hacmp.out**, which keeps a record of all of the local cluster events as performed by the HACMP event scripts. These HACMP event scripts automate many common system administration tasks, and, in the event of a failure, will manage HACMP and system resource to provide recovery.
- **Recover From HACMP Script Failure.** Contains a command that HACMP will run to recover from a script failure. This is useful if the Cluster Manager is in reconfiguration due to a failed event script. Use this option after having manually fixed the error condition.
- **Restore HACMP Configuration Database from Active Configuration.**

Viewing Cluster Information (clstat) in WebSMIT

With HACMP 5.4.1, you can use WebSMIT to:

- Display detailed cluster information
- Navigate and view the status of the running cluster
- Configure and manage the cluster
- View graphical displays of sites, networks, nodes and resource group dependencies.

Useful Commands

You have these additional utilities:

- To view the resource group location and status, use the **clRGinfo** command.
- To view the service IP label information, run the **ifconfig** command on the node that currently owns the resource group.

For a list of commands supported in HACMP for Linux, see [Command Reference](#) in [Appendix A: Command Reference and the clinfo Utility](#).

Logging Messages

HACMP for Linux uses the standard logging facilities for HACMP. For information about logging in HACMP, see the HACMP for AIX *Troubleshooting Guide*.

To troubleshoot the HACMP operations in your cluster, use the event summaries in the **hacmp.out** file and **syslog**.

The system logs messages into the following files:

- **/tmp/clstrmgr.debug**
- **/tmp/cspoc.log**
- **/tmp/clappmond**
- **/tmp/hacmp.out**
- **/usr/es/adm/cluster.log**
- **/var/hacmp/clcomd/clcomd.log**
- **/var/hacmp/clcomd/clcomddiag.log**
- **/var/hacmp/log/clutils.log**
- **/usr/es/sbin/cluster/wsm/logs/wsm_smit.***
- **<APACHE_HOME>/websmit/logs/wsm_smit.***
- **/usr/es/sbin/cluster/snapshots/<snapshot_name>***

Collecting Cluster Log Files for Problem Reporting

To view the system files and log files as they are collected in an archive file:

1. In WebSMIT, go to the **Collect Cluster log files for Problem Reporting** menu.
2. Type or select values in entry fields.
3. Use an appropriate Linux tool to extract or view the archive file. The archive file contains the log and system files.

Solving Common Problems with Networks and Applications

This section contains the following topics:

- [Identifying Causes of Unexpected Network and Network Interface Failures](#)
- [Troubleshooting an Unsuccessful Application Fallover to Another Node](#)
- [Troubleshooting the Serial Connection.](#)

For a list of other specific problems and tips on how to check for them, see the HACMP for AIX *Troubleshooting Guide*.

Identifying Causes of Unexpected Network and Network Interface Failures

This section lists some of the possible causes of the network errors you may receive.

For example, HACMP logs **network_down** or **interface_failed** events, but the NIC appears to be functional.

Verify that the information defined in the HACMP configuration matches the information displayed by `ifconfig`, for example:

To verify the HACMP configuration for node `ppstest2`, run the `cllsif` command:

```
ppstest2:~ # /usr/es/sbin/cluster/utilities/cllsif -c | grep ppstest2
ppstest2_enstby1:boot:net_ether_01:ether:public:ppstest2:192.9.201.131::
eth1::255.255.255.128::

ppstest2_enboot:boot:net_ether_01:ether:public:ppstest2:192.9.201.2::eth
0::255.255.255.128::

ppstest2:boot:net_token_01:token:public:ppstest2:9.57.28.4::tr0::255.255
.255.128::
```

Now run the `ifconfig` command on node `ppstest2` and compare the results:

```
ppstest2:~ # ifconfig
eth0      Link encap:Ethernet  HWaddr 00:06:29:DC:82:7A
          inet addr:192.9.201.2  Bcast:192.9.201.127
          Mask:255.255.255.128
          inet6 addr: fe80::206:29ff:fedc:827a/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:7120 errors:0 dropped:0 overruns:0 frame:0
          TX packets:12605 errors:0 dropped:0 overruns:0 carrier:0
          collisions:3821 txqueuelen:1000
          RX bytes:499630 (487.9 Kb)  TX bytes:15802477 (15.0 Mb)
          Interrupt:37 Base address:0xec00

eth1      Link encap:Ethernet  HWaddr 00:06:29:DC:E0:2A
          inet addr:192.9.201.131  Bcast:192.9.201.255
          Mask:255.255.255.128
          inet6 addr: fe80::206:29ff:fedc:e02a/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:10695 errors:0 dropped:0 overruns:0 frame:0
          TX packets:14381 errors:0 dropped:0 overruns:0 carrier:0
          collisions:4205 txqueuelen:1000
          RX bytes:5690673 (5.4 Mb)  TX bytes:15921948 (15.1 Mb)
          Interrupt:38 Base address:0xec00
```

```

lo          Link encap:Local Loopback
            inet addr:127.0.0.1  Mask:255.0.0.0
            inet6 addr: ::1/128 Scope:Host
            UP LOOPBACK RUNNING  MTU:16436  Metric:1
            RX packets:102 errors:0 dropped:0 overruns:0 frame:0
            TX packets:102 errors:0 dropped:0 overruns:0 carrier:0
            collisions:0 txqueuelen:0
            RX bytes:6624 (6.4 Kb)  TX bytes:6624 (6.4 Kb)

tr0        Link encap:16/4 Mbps Token Ring (New)  HWaddr
00:60:94:8A:D2:F7
            inet addr:9.57.28.4  Bcast:9.57.28.127  Mask:255.255.255.128
            inet6 addr: fe80::260:94ff:fe8a:d2f7/64 Scope:Link
            UP BROADCAST RUNNING MULTICAST  MTU:4056  Metric:1
            RX packets:2332 errors:0 dropped:0 overruns:0 frame:0
            TX packets:2577 errors:0 dropped:0 overruns:0 carrier:0
            collisions:0 txqueuelen:100
            RX bytes:431872 (421.7 Kb)  TX bytes:829909 (810.4 Kb)
            Interrupt:52 Base address:0xec00
  
```

Notice that the interface name and IP address for each address defined to HACMP (lo) is never defined to HACMP.

To fix the problem, do either of the following:

- If the interface name or IP address is incorrect, use the **Change/Show Communication Interfaces/Devices** WebSMIT menu to change them.
- If the netmask is incorrect, use the **Change/Show Networks** menu to change the netmask for all interfaces on this network.

Note: You do *not* have to specify interface names for service IP addresses or for persistent addresses. HACMP keeps these addresses available by moving them to a different interface after any failure.

Troubleshooting an Unsuccessful Application Fallover to Another Node

If your application starts successfully on one node, but HACMP issues an EVENT FAILED message when trying to perform an application fallover to another node, make sure that the application start, stop and notify scripts exist and are executable on every node in the cluster. Use the `cllsserv` command.

For example:

```

ppstest2:~ # /usr/es/sbin/cluster/utilities/cllsserv
app_test2_primary /usr/local/app_start /usr/local/app_stop
ppstest2:~ # ls -l /usr/local/app_start
-rwxr--r--  1 root root 169 May 10 22:54 /usr/local/app_start
  
```

Troubleshooting the Serial Connection

This section describes how to test an installed serial connection.

To ensure that the RS232 cable is properly configured and transmits data, run the following test after creating the `tty` device on both nodes.

Run this test while the `tty` device is *not* in use. If the cluster is active, remove the serial network dynamically from the configuration before running the test. Also, verify that the `tty` device is *not* in use by any other process.

To determine if the device is in use, run the **fuser** command:

```
fuser /dev/tty0
```

The output lists the PID of any process that uses the device.

If the device is in use by RSCT, the output shows that a process `hats_rs232_nim` is accessing the device. After the network has been dynamically removed from the cluster configuration, no such process should exist.

In rare cases, the `hats_rs232_nim` process may *not* terminate during a dynamic removal of the network or a stop of the cluster services. In these cases, you should call IBM support. However, it is safe to terminate any leftover `hats_nim_rs232` process if the cluster is inactive on the local node.

Use the **fuser** command to terminate a process that accesses the **tty** device:

```
fuser -k /dev/tty0
```

Running the sttyTest

The **stty** test determines whether the serial connection allows the transmission of communications.

Running the stty Test on TTYs with RTS Flow Control Set

To perform the **stty** test:

1. On the receiving side, run:

```
(stty raw -echo; cat > outputfilename) < /dev/tty2
```

2. On the sending side, run:

```
(stty raw -echo < /dev/tty1; cat filetobesent ; sleep 5) > /dev/tty1
```

Running the stty Test on TTY's with XON or No Flow Control Set:

To perform the **stty** test:

1. On the receiving side (node 2), run:

```
(stty raw -echo ixon ixoff; cat > outputfilename) < /dev/tty2
```

2. On the sending side, run

```
(stty raw -echo ixon ixoff < /dev/tty1; cat filetobesent; sleep 5) > /dev/tty1
```

If the nodes are able to communicate over the serial cable, both nodes display their **tty** settings and return to the prompt.

If the data is transmitted successfully from one node to another, then the text from the `/etc/hosts` file from the second node appears on the console of the first node. Note that you can use any text file for this test, and do *not* need to specifically use the `/etc/hosts` file.

After you install and test the serial connection, you define the connection as a point-to-point network to HACMP. For information about how to configure a serial network, see [Configuring Serial Networks for Heartbeating](#) in [Chapter 3: Common Task Summary](#).

5

Monitoring and Troubleshooting a Cluster

Solving Common Problems with Networks and Applications

Appendix A: Command Reference and the clinfo Utility

This appendix includes:

- [Command Reference](#)
- [The clinfo Utility: Providing APIs to Interact with Other Applications.](#)

Command Reference

This section lists HACMP commands that are supported on HACMP for Linux. For further information on the commands, use their man pages. The command directory is `/usr/es/sbin/cluster/man/cat1`.

The majority of the utilities packaged with HACMP for AIX are packaged with HACMP for Linux as well. However, some of the utilities may *not* fully work in release 5.4.1.

Note: The list of commands provided in this section is *not* guaranteed to be up-to-date or exhaustive, and is provided here as a reference only. In general, the HACMP user interface lets you perform all necessary actions in it, without having to use HACMP commands directly. The general rule of thumb is that everything you need to do is available to you via options in WebSMIT. If, however, a particular user interface option is *not* available in HACMP for Linux (it may be available in HACMP on AIX), the corresponding CLI is also *not* supported on the Linux platform.

clstop	clgetaddr	clcallev
clruncmd	cllgrp	clcomd_ctrl
cllscf	cllsparm	clcheck_server
clstat	get_local_nodename	cltopinfo
clshowsrv	cllsnim	clRGinfo
cllres	clsnapshot	clRMupdate
cllsserv	clsnapshotinfo	clRGmove
clshowres	clgetactivenode	
cllsev	clfindres	

The **clinfo** Utility: Providing APIs to Interact with Other Applications

The HACMP software extends the benefits of highly available servers, data, and applications to client software applications by providing the clients with notification of cluster state changes through the Cluster Manager and the Cluster Information Program (**clinfo**).

The **clinfo** utility and its associated APIs enable developers to write applications that recognize and respond to changes in a cluster.

The **clinfo** utility includes:

- Its own API. The **clinfo** utility provides the client API library. Use this library in client applications that you create. Your client applications can use socket connections to **clinfo** to retrieve cluster information.

The **clinfo** API functions provide application developers with both a C and a C++ language interface for accessing cluster status information. The HACMP software includes two versions of the **clinfo** APIs: one for single-threaded applications and one for multi-threaded applications.

- Scripts to communicate with the Cluster Manager. The **clinfo** utility gets cluster information from the Cluster Manager and allows clients communicating with it to be aware of a cluster's state changes.

Note: HACMP for Linux does *not* support SNMP as its communications protocol and therefore **clinfo** implementation in HACMP for Linux is different from **clinfo** implementation in HACMP for AIX. On the Linux platform, the Cluster Manager services the **clinfo** client library directly.

At a high level, the **clinfo** utility works as follows:

- The Cluster Manager, **clstrmgr**, listens on the **clinfo** client socket for **clinfo** API requests.
- The **clinfo** utility takes the cluster configuration information from the **clstrmgr** internal data structures.
- The **clinfo** utility stores information (retrieved from the Cluster Manager) in series of list structures, with a list for cluster nodes, networks and other cluster components.
- The client library API reads from the list structures.

For complete information on **clinfo** and clients library API, see the HACMP for AIX *Programming Clients Applications Guide*.

Glossary

This Glossary contains definitions of terms used in the *HACMP for Linux: Installation and Administration Guide*. The definitions assume a general familiarity with terminology for HACMP™ and on Linux.

Also see the *HACMP Master Glossary*.

A

application

A service, such as a database, or a collection of system services and their dependent resources, such as a service IP label associated with your application, or an application server, that you want to keep highly available with the use of HACMP.

application server

A collection of an application start and stop scripts that you provide to HACMP, by entering the pathnames for the scripts in the WebSMIT user interface. An application server is “a resource” associated with an application, you include it in a *resource group* and HACMP ensures its high availability (the fact that the application can start and stop successfully no matter on which cluster node it is currently running).

C

cluster IP network

A cluster IP network is used for cluster communications between the nodes and for sending heartbeating information. All IP labels configured on the same HACMP network share the netmask, but may be required to have different subnets. You also must add the IP addresses for an IP cluster network to the */etc/hosts* file on each node.

cluster node

A physical machine, typically an AIX or a Linux server on which you install HACMP. A cluster node also hosts an application. A cluster

node serves as a server for application’s clients. HACMP’s role is to ensure continuous access to the application, no matter on which node in the cluster it is currently active.

See also *home node* and *takeover node*.

cluster resources

Resources that include volume groups, an application server, and a service IP label. All or some of these resources can be associated with an application you plan to keep highly available. You include cluster resources into *resource groups*.

cluster startup

The starting of HACMP cluster services on the node(s).

cluster shutdown

The stopping of HACMP cluster services on the node(s).

F

fallback

An action of a resource group, when HACMP returns it from a *takeover node* back to the *home node*. You choose a *fallback resource group policy* from a predefined list in WebSMIT.

See also *resource group startup* and *fallover*.

fallover

An action of a resource group, when HACMP moves it from one node to another. In other words, a resource group and its associated application *fall over* to another node. You choose a *resource group fallover policy* from a predefined list in WebSMIT.

See also *resource group startup*, and *fallback*.

H

H

heartbeat

State-of-health message exchanged between network modules. HACMP uses these to track the membership and status of the cluster nodes. A means of detecting failure in the cluster.

home node

A node on which the application is hosted based on your default configuration, and under normal conditions.

I

IP label

A name of a network interface card (NIC) that you provide to HACMP. Network configuration for HACMP requires planning for several types of IP labels:

- base (or boot) IP labels
- service IP labels (the ones through which a connection for a highly available application is established)
- (optional) backup IP labels.
- IP aliases for HACMP network configuration.
- Node-bound persistent IP labels on each node.

To ensure high availability and access to the application, HACMP “moves” the service IP label/address associated with the application to another node in the cluster, using the function called IPAT via IP Aliasing.

See also *IP Address Takeover via IP Aliasing (IPAT via IP Aliasing)*.

IP alias

An alias placed on an IP label. It coexists on an interface along with the IP label. Networks that support Gratuitous ARP cache updates enable configuration of IP aliases.

See also *IP Address Takeover via IP Aliasing (IPAT via IP Aliasing)*.

IP Address Takeover (IPAT)

A process whereby a service IP label on one node is taken over (or moved) to a backup node in the cluster. HACMP uses IPAT to provide

high availability of IP service labels that belong to resource groups (and provide access to applications).

HACMP for Linux supports IPAT via IP Aliasing.

IP Address Takeover via IP Aliasing

A default method of IPAT used in HACMP for AIX and also the only method of IPAT used in HACMP for Linux.

HACMP uses IPAT via IP Aliasing in cases when it must automatically recover a service IP label on another node. To configure IPAT via IP Aliasing, you configure service IP labels and their aliases to the system. When HACMP performs IPAT during automatic cluster events, it places an IP alias (recovered from the “failed” node) on top of the service IP address on the takeover node. Access to the application is preserved.

P

partitioned cluster

Complete loss of communication between nodes in the cluster. Occurs when all TCP/IP-based networks connecting cluster nodes fail, if a global cluster network is *not* defined. One or more nodes are completely isolated from the other nodes.

pre- and post-events

Customized scripts provided by you (or other system administrators), that you can make known to HACMP and that will be run prior a particular cluster event, or after a particular cluster event.

For more information on pre- and post-event scripts, see the chapter on *Planning Cluster Events* in the *HACMP for AIX Planning Guide*.

R

resource group

A collection of cluster resources. *Cluster resources* can include volume groups, an application server, and a service IP label. All or some of these resources can be associated with an application you plan to keep highly available.

See also *cluster resources*.

resource group startup

An activation of a resource group and its associated resources on a specified cluster node. You choose a resource group startup policy from a predefined list in WebSMIT.

resource group fallback

See *fallback*.

resource group fallover

See *fallover*.

resource group takeover

See *takeover*.

RSCT (Reliable Scalable Cluster Technology)

IBM RSCT services provide scalability, notify distributed subsystems of software failure, and coordinate recovery and synchronization among all subsystems in the software stack. This package of services is a prerequisite for HACMP for Linux.

takeover node

A backup cluster node on which HACMP may move the application, either on your request (for instance, if you want to perform planned maintenance on the home node), or automatically, due to a cluster component failure. In HACMP for Linux v.5.4.1, a cluster configuration includes up to eight nodes, therefore, you can have more than one takeover node for a particular application.

S
startup

see *cluster startup*.

selective fallover

An automatically launched function of HACMP that attempts to selectively move only a resource group that has been affected by an individual resource failure to another node in the cluster, rather than moving all resource groups.

shutdown

See *cluster shutdown*.

synchronization

An HACMP function during which the system ensures that the user-configured information is synchronized across all currently active nodes in the cluster.

T
takeover

An action during which HACMP takes over resources from one node and moves them to another node. A backup node is referred to as a *takeover node*.

T

Notices for HACMP for Linux: Installation and Administration Guide

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

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

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

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

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

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

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

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

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

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

IBM Corporation
Dept. LRAS / Bldg. 003
11400 Burnet Road
Austin, TX 78758-3493
U.S.A.

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

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

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

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

Index

+.*/-

- 28, 33
- /etc/hosts 27, 36
- /tmp/hacmp.out 19, 64
- /usr/es/sbin/cluster/clstat 52
- /usr/es/sbin/cluster/etc/inittab 33
- /usr/es/sbin/cluster/install_hacmp 30
- /usr/es/sbin/cluster/man/cat1
 - command directory 69
- /usr/es/sbin/cluster/netmon.cf file 57
- /usr/es/sbin/cluster/release_notes.linux 31
- /usr/es/sbin/cluster/snapshots/* 64
- /usr/es/sbin/cluster/utilities/clsif -c 65
- /usr/es/sbin/cluster/utilities/clserv 50, 66
- /usr/es/sbin/cluster/wsm/logs/wsm_smit.* 64
- /usr/es/sbin/cluster/wsm/README file
 - WebSMIT information 31, 37
- /var/hacmp/log/hacmp.install.log
 - installation log 32

A

- adding
 - IP label and application server to a resource group 51
 - predefined communication interface/device 43
- Apache server
 - integration with WebSMIT 32
- Apache WebServer 28
- API
 - for interaction with other software clients 70
- application server
 - checking the scripts on each node 66
- application servers
 - configuring with extended path 49
 - start script 49
 - stop script 49
- applications
 - managing in HACMP 54

B

- boot interfaces
 - configuring to HACMP 44

C

- changing resources 51
- clcomd 33
- clfindres 69
- clRGinfo 69

- clRGmove 69
- clsnapshot 69
- clsnapshot.log file 60
- clstat 69
 - output in ASCII format 53
 - viewing cluster status 52
- clstop 69
- cltopinfo 69
- Cluster Communications utility, clcomd 33
- cluster events
 - resetting customizations 60
- cluster name
 - adding 39
 - changing 40
- cluster network
 - definition 15
- cluster partitioning
 - preventing 19
- cluster resources
 - adding to resource groups 51
 - definition 16
- Cluster Single Point of Control (C-SPOC) 54
- cluster snapshot
 - creating prior to resetting tunable values 60
- cluster status
 - clstat utility 52
 - obtaining 53
 - viewing 52
- cluster topology
 - configuring with extended path 39
- commands
 - supported in HACMP for Linux 69
- communication devices on serial networks
 - adding to the HACMP cluster 44
- communication interface/device
 - configuring 42, 43
- communication paths
 - establishing for cluster 40
- communications interface
 - assigning persistent node IP labels 44
- configuration
 - prerequisites 36
- configuring
 - application servers 49
 - cluster topology with extended path 39
 - communication interfaces/devices 42
 - HACMP networks 36
 - heartbeat paths 36
 - persistent node IP label 45

D

- DARE 36
- displaying
 - applications, networks, nodes, resource groups 52
- displaying cluster information
 - WebSMIT 63
- dynamic reconfiguration
 - DARE 25
- Dynamic reconfiguration of cluster resources (DARE) 36

F

- failure of the network interface in HACMP 19
- file systems
 - not supported as resources 22
- files
 - /usr/es/sbin/cluster/netmon.cf 57
 - netmon 57
- filesets
 - hacmp.client 29
 - hacmp.server 29
- filesets for installation 29
- fuser command 67

H

- HACMP cluster services
 - starting 54
 - stopping 54
- HACMP for Linux
 - software supported 28
- HACMP logs
 - list 64
- hacmp.client
 - filesets 29
- hacmp.out file
 - HACMP logs 63
- hacmp.server
 - filesets 29
- hardware
 - for cluster nodes 21
- heartbeat paths
 - configuring 36, 37
- heartbeating addresses
 - IP address offset 42
- heartbeating network 26

I

- ifconfig
 - output 46
- images for installation 29
- installation log file name 32
- installation script 30

- installing HACMP for Linux
 - filesets 29
 - images 29
 - installation media 29
 - software supported 28
- IP address
 - offset for heartbeating over IP Aliases 42
- IP label/address
 - resource 48
- IP labels
 - recovery mechanism 16
- IPAT via IP aliasing 16
- IPSEC 34

K

- ksh93 28

L

- license agreement
 - accepting during installation 30
- limitations 36
- list of log files 64
- listing application servers 50, 66
- listing HACMP network interfaces 65
- log files
 - collecting 64
- log viewing
 - where to locate in WebSMIT 54
- logging 64
- logging facilities 64
- logs
 - Apache 64
 - hacmp.out file 63
 - viewing and managing 55
- lssrc 33

M

- man pages, directory and list 69
- managing
 - logs 55

N

- naming
 - cluster 40
- netmasks
 - troubleshooting 66
- netmon configuration file 57
- network interface failure
 - how HACMP handles it 19
- network interfaces
 - configuring predefined to cluster 43
 - listing defined 65

- networks
 - configuring using extended path 36
 - single adapter configuration 57
- NFS
 - not supported as cluster resources 22
- nodes in HACMP cluster
 - hardware 21
- number of nodes 21

P

- performance 24
- persistent IP-labels
 - changing and deleting 46
- persistent node IP label
 - assigning 44
 - configuring 45
 - removing 45
- planning
 - cluster nodes 24
 - resource groups 25
 - sites 24
- Pluggable Authentication Module package (PAM)
 - security for WebSMIT 34
- port for the HACMP for Linux software
 - access online 37
- prerequisites, prior to configuring 36
- propagating information from one node to all
 - synchronizing a cluster 51

R

- recovery actions
 - for a local network failure event 19
- release notes 31
- removing
 - persistent node IP label 45
- reserved words 56
- resetting cluster tunables
 - how it works 60
- resource groups
 - definition 16
 - moving to other nodes, taking online and offline 54
 - non-concurrent 25
 - planning for, in HACMP for Linux environment 25
 - startup, fallover, fallback policies 26
- resources in a resource group
 - changing 51
- RHEL 28
- rpm command 31
- RS232 serial lines
 - testing 67
- RSCT 28

S

- sample two-node cluster
 - diagram 17
- security
 - clcomd 33
 - PAM 33
- security for WebSMIT
 - Pluggable Authentication Module (PAM) 34
- serial devices
 - configuring to HACMP 44
- serial networks
 - testing
 - RS232 line 67
 - troubleshooting 67
- service IP label
 - definition 15
- single-adapter networks and netmon configuration file 57
- sites 25
- SLES 28
- snapshot
 - log file 64
- software versions
 - installation 28
- starting
 - HACMP cluster services 54
- stopping HACMP cluster services 54
- synchronizing a cluster 51
- system directories
 - entries added post-installation 33
- System Management (C-SPOC) 37

T

- tasks to do before configuring a cluster 36
- testing
 - serial connection 67
- troubleshooting
 - viewing and managing logs 55
- troubleshooting a serial network connection 67

U

- user interface for configuration, WebSMIT
 - installing 31

V

- viewing
 - the cluster status 52
- viewing configured
 - applications, networks, nodes, resource groups 52

Index
W – W

W

WebSMIT

cluster management 31

display cluster information 63

installing 31