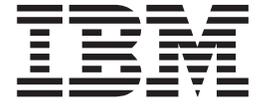


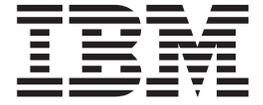
IBM Cluster Systems Management for Linux[®]



Set-Up HOWTO

Version 1 Release 1

IBM Cluster Systems Management for Linux[®]



Set-Up HOWTO

Version 1 Release 1

Note!

Before using this information and the product it supports, read the information in "Appendix B. Notices" on page 15.

First Edition (June 2001)

This edition of the *IBM Cluster Systems Management for Linux Set-Up HOWTO* applies to IBM Cluster Systems Management for Linux Version 1 Release 1, program number 5799-GNJ, and to all subsequent releases of this product until otherwise indicated in new editions.

IBM welcomes your comments. A form for readers' comments may be provided at the back of this publication, or you may address your comments to the following address:

International Business Machines Corporation
Department 55JA, Mail Station P384
2455 South Road
Poughkeepsie, NY 12601-5400
United States of America

FAX (United States & Canada): 1+845+432-9405

FAX (Other Countries):

Your International Access Code +1+845+432-9405

IBMLink™ (United States customers only): IBMUSM10(MHVRCFS)

IBM Mail Exchange: USIB6TC9 at IBMMAIL

Internet e-mail: mhvrdfs@us.ibm.com

If you would like a reply, be sure to include your name, address, telephone number, or FAX number.

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.

© **Copyright International Business Machines Corporation 2001. All rights reserved.**

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

Contents

About This HOWTO	v
Who Should Use This HOWTO	v
Typographic Conventions	v
Related Information	v
How to Obtain Publications	vi
Setting Up IBM Cluster Systems Management for Linux	1
Specified Operating Environment	1
Hardware Requirements	1
Software Requirements	1
Planning for Cluster Systems Management	2
Pre-installation tasks	2
Overview of Cluster Systems Management installation	4
Running installms	4
Running definenode or addnode	5
Prerequisites for running definenode or addnode	5
Run definenode or addnode	5
Example of definenode command run interactively	6
Verifying node definitions	7
Changing internal service processor user IDs and passwords	7
Running installnode	7
Prerequisites for running installnode	7
Run installnode	8
Adding a node to an existing cluster	8
Removing a node from an existing cluster	8
Getting started with the newly installed cluster	9
FAQs, hints, tips, and troubleshooting	9
ACL File FAQs	10
Appendix A. Node Attributes Template	13
Appendix B. Notices	15
Trademarks	16
Publicly Available Software	16
Index	19

About This HOWTO

This HOWTO provides the specified operating environment for the IBM Cluster Systems Management for Linux (CSM) set of tools and describes how to install and set up a CSM cluster on an existing group of nodes.

Who Should Use This HOWTO

This HOWTO is intended for system administrators who want to use IBM Cluster Systems Management for Linux. It describes tools that are provided to make the installation of Cluster Systems Management easier. The system administrator should have experience in UNIX[®] administration and networked systems.

Typographic Conventions

This book uses the following typographic conventions:

Typographic	Usage
Bold	<ul style="list-style-type: none">• Bold words or characters represent system elements that you must use literally, such as commands, flags, and path names.
<i>Italic</i>	<ul style="list-style-type: none">• <i>Italic</i> words or characters represent variable values that you must supply.• <i>Italics</i> are also used for book titles and for general emphasis in text.
Constant width	Examples and information that the system displays appear in constant width typeface.
[]	Brackets enclose optional items in format and syntax descriptions.
{ }	Braces enclose a list from which you must choose an item in format and syntax descriptions.
	A vertical bar separates items in a list of choices. (In other words, it means “or.”)
< >	Angle brackets (less-than and greater-than) enclose the name of a key on the keyboard. For example, <Enter> refers to the key on your terminal or workstation that is labeled with the word Enter.
...	An ellipsis indicates that you can repeat the preceding item one or more times.
<Ctrl-x>	The notation <Ctrl-x> indicates a control character sequence. For example, <Ctrl-c> means that you hold down the control key while pressing <c> .
\	The continuation character is used in coding examples in this book for formatting purposes.

Related Information

IBM Cluster Systems Management for Linux Monitoring HOWTO, SA22-7852-00

IBM Cluster Systems Management for Linux Overview HOWTO, SA22-7857-00

IBM Cluster Systems Management for Linux Remote Control HOWTO, SA22-7856-00

IBM Cluster Systems Management for Linux Technical Reference, SA22-7851-00

How to Obtain Publications

The IBM Cluster Systems Management for Linux publications are available as HTML and PDF files on the CD-ROM in the **/doc** directory or on the installed system in the **/opt/csm/doc** directory.

README information is available on the CD-ROM in the root directory (*/*).

The file names are as follows:

- *IBM Cluster Systems Management for Linux Monitoring HOWTO*, csmadm.pdf
- *IBM Cluster Systems Management for Linux Remote Control HOWTO*, csmremot.pdf
- *IBM Cluster Systems Management for Linux Set-Up HOWTO*, csmsetup.pdf
- *IBM Cluster Systems Management for Linux Overview HOWTO*, csmovrvw.pdf
- *IBM Cluster Systems Management for Linux Technical Reference*, csmtech.pdf

Publications for IBM Cluster Systems Management for Linux were available also at the time of this release at the following URL:

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

Setting Up IBM Cluster Systems Management for Linux

The set-up process helps a system administrator get IBM Cluster Systems Management for Linux (CSM), hereafter known as Cluster Systems Management, up and running easily by setting up a management server and managed nodes on existing Linux systems. The application is available as an installation image in a directory or on the CSM CD-ROM. This document describes the minimum hardware and software requirements needed to use this product. See the “Specified Operating Environment”.

Information is also provided about planning and pre-installation tasks that you need to perform so that the installation will go smoothly and easily. Next, there is a step-by-step procedure for installing and setting up the cluster on an existing group of nodes. Finally, a troubleshooting section is provided in the form of frequently asked questions. You should read this document carefully and be familiar with it throughout before beginning the installation and set-up tasks.

Specified Operating Environment

This section describes the hardware and software that are required for IBM Cluster Systems Management for Linux. For more detailed information, see the announcement.

Hardware Requirements

This product is supported on the IBM xSeries 330 and 340.

Remote Control Hardware Requirements

To support remote control, the following hardware is required:

- One ASM PCI adapter for every 10 nodes.
- An Internal Service Processor per node
- Equinox Serial Provider 8 or Equinox Serial Provider 16
- An Ethernet adapter for each virtual LAN (VLAN)
- Appropriate cabling

Note: The total number of nodes should be less than or equal to the number of ports associated with all of the Equinox Serial Providers that are installed.

For more details, see *IBM Cluster Systems Management for Linux Remote Control HOWTO*.

Software Requirements

IBM Cluster Systems Management for Linux has requirements for non-IBM software as well as IBM-developed software. As a convenience, the required software that is not part of the Red Hat distribution is included on the CSM CD-ROM. Unless otherwise specified, the software is required on the management server and on the managed nodes. The non-IBM software required is as follows:

1. Red Hat Linux 7.1 distribution
2. glibc 2.2
3. libstdc++ 2.9
4. pdksh 5.0
5. perl 5.00503
6. make 3.7
7. nfs-utils 0.1 (This is on the management server only.)
8. DBI 1.14
9. Text-CSV_XS 0.21
10. SQL-Statement 0.1016
11. DBD-CSV 0.1024

12. Cfengine 1.6.2
13. Pidentd 3.0

The following IBM software packages are required:

1. csm.client (This is required on the managed nodes only.)
2. csm.core
3. csm.server (This is required on the management server only.)
4. csm.dsh
5. rsct.core
6. rsct.core.utils
7. src

Network Requirements

TCP/IP and a network adapter

Network Considerations: In configuring a Cluster Systems Management cluster, give particular attention to the following:

- To secure remote control functions, IBM suggests that a management server be defined and isolated on a private network.

Memory and Disk Space

For the management server, a minimum of 128MB of memory and 120MB of disk space is required.

For the managed node, a minimum of 128MB of memory and 20MB of disk space is required.

Limitations

A cluster of up to 32 nodes is supported.

Note: The minimum configuration and sizing of each node is highly dependent on the user's application and performance requirements.

Planning for Cluster Systems Management

See *IBM Cluster Systems Management for Linux Remote Control HOWTO* for information on hardware configuration considerations and set-up instructions, including a filled out node-attribute table that you can use for guidance in filling out your own blank template. A blank template is provided in "Appendix A. Node Attributes Template" on page 13.

Pre-installation tasks

There are several tasks that the administrator must do to prepare for installation of Cluster Systems Management:

1. Information needs to be gathered to fill out the template in "Appendix A. Node Attributes Template" on page 13. See the section "Hardware Configuration" in *IBM Cluster Systems Management for Linux Remote Control HOWTO* for a detailed discussion of hardware configuration and set-up considerations and for an example of a filled-out node-attribute table.
2. The host names of the nodes that are being defined to the cluster need to be registered with the nameserver or added to the **/etc/hosts** file on the management server.
3. IBM® suggests that you create a separate partition called **/fttpboot** on the management server to contain a copy of the packages on the installation media before you install CSM. See "Creating the /fttpboot partition" on page 3 for the procedure that shows how to do this.
4. IBM suggests that you add **/opt/bin** to your \$PATH variable and **/opt/man** to your \$MANPATH variable.

Creating the /tftpboot partition

Before installing CSM, create a partition called **/tftpboot** that consists of 100MB of space. This will hold the required RPM and tarball packages for installation.

To create **/tftpboot** by using **fdisk**, do the following:

1. Log in as root.
2. Run **fdisk**.
3. Move the cursor down to the **Free Space** entry.
4. Press **n** to create a new partition
5. Enter 100 for the Size (in MB).
6. Select **Beginning** to specify where to position the partition.
7. With the new partition selected, press **t** to select the partition type.
8. Enter "83" (Linux) as the partition type.
9. Select **Write** to write the new partition information to disk.

Make a note of the device name and number of the new partition because you will need it for the next steps. Examples of the new partition name might be similar to **/dev/hda7** or **/dev/sda8**.

10. Reboot the system after the partition has been created. Then do the following:

- a. Run the following command:

```
mkfs /dev/device
```

(where *device* is the name of the new partition; for example, **hda7** or **sda8**.)

- b. Then run:

```
mkdir -p /tftpboot
```

- c. Then mount the partition:

```
mount /dev/device /tftpboot
```

- d. Add the following line to **/etc/fstab**:

```
/dev/device /tftpboot ext2 defaults 1 2
```

Preparing for dsh and configuring the remote shell

A distributed shell program (**dsh**) is used to run commands on the nodes. It is contained in the **csmdsh** RPM and installed by the **installms** command. The **dsh** program uses a remote shell of your choice to issue remote commands to the managed nodes from the management server. The following preparation to enable the remote shell is required on each node before **dsh** is installed:

1. Decide on the remote shell that you will be using. The default remote shell is **rsh**.
2. If you are using **rsh**, make sure the rsh-server RPM is installed and running. If you are using another remote shell, make sure its daemon is installed and running.
3. Use the fully qualified host name when you define a node for the remote shell. If the remote shell requires a list of nodes in its configuration, then the nodes must be defined by their fully qualified host names. This allows the **dsh** command to recognize the node. You can also use an alias to define a node. Aliases are permitted provided the fully qualified host name is also given.
4. Security must be set up on each node in such a way that **dsh** is allowed to run commands on that node. If you are using **rsh**, add the management server host name to the **/root/.rhosts** file on the nodes that will be managed nodes. If you are using another shell, you must fulfill the requirements for installing and using that shell.

The **DSH_REMOTE_CMD** environment variable is used to specify a remote shell other than the default. This environment variable should always be set when CSM commands are issued because some CSM commands use **dsh** internally and will use **rsh** as the default if **DSH_REMOTE_CMD** is not set.

5. To ensure that **dsh** is working on each of the nodes, use the remote shell to run a remote command from the node that will be the management server to each node that will be in the cluster.

Note: Be aware that the **dsh** command does not provide the requirements for a specific security configuration. The user is responsible for fulfilling the particular security obligations of a specified security environment. At a minimum, you can configure **rsh** with the **/root/.rhosts** file on nodes. A more secure environment might have Kerberos configured or might be using some type of shell that conforms to the IETF (Internet Engineering Task Force) Secure Shell protocol.

For more information on **dsh**, see the man page or the *IBM Cluster Systems Management for Linux Technical Reference*.

Overview of Cluster Systems Management installation

To install Cluster Systems Management, you can take one of two approaches. For a simple installation without interim verification, you can run the following commands in sequence:

1. Run **installms**.
2. Run **addnode**.

The **addnode** command runs **definenode** and then **installnode** automatically.

If you need more control and would like the ability to doublecheck and make interim changes during installation, then run **definenode** and then **installnode** as follows:

1. Run the **installms** command.
2. Run the **definenode** command.
 - a. Use **lsnode -P** to view node definitions.
 - b. Use **chnode -P** to make changes to the attributes of a node.
3. Run the **installnode** command.

All of these commands are run on the management server. Details on these commands can be found in their man pages or in *IBM Cluster Systems Management for Linux Technical Reference*.

Attention:

After each node is installed by running **installnode**, you need to reboot the node to enable remote console support. The node does not have to be rebooted if remote console support is not being used.

The **identd** service is started for you on the management server and the managed nodes by the installation process. See the *Security* section of the *IBM Cluster Systems Management for Linux Overview HOWTO* for more information about **identd**.

Running installms

The **installms** command performs the tasks that are necessary to make this system a management server. It installs the appropriate software listed in “Specified Operating Environment” on page 1 on the management server automatically if it is not already installed or if it is installed at a previous level.

IBM suggests that you set up the **/tftpboot** partition before you run **installms**. You also need to provide and mount the CSM distribution CD-ROM. The default mount point is **/mnt/cdrom**.

1. Log in as root to the machine that is to become the management server.
2. Copy the **installms** program from the CD-ROM to a temporary directory. The **/tmp** directory is fine. The program first copies installation packages from a download directory or from the CD-ROM that contains the CSM application to **/tftpboot/rpm** and **/tftpboot/tarball**.
3. Next, the program asks you to insert the Linux distribution CD-ROMs in the drive. The program automatically mounts and unmounts the CD-ROMs as needed.

4. The RPM, CSM, and tarball packages are used to install the code and initialize the server. The **installms** command determines whether the code needs to be installed (if it is missing or backlevel) and then installs or replaces the required packages, as necessary.

Note: See the README for the latest information on obtaining the required packages for CSM. For example, you may need to download some tarball packages from the World-Wide Web.

After **installms** has been run successfully, run **definnode** or **addnode**.

For more information on **installms**, see the man page or the *IBM Cluster Systems Management for Linux Technical Reference*.

Running definnode or addnode

After the management server is installed by running **installms**, run **definnode** to define all of the nodes in the cluster or **addnode** to define and install the nodes in the cluster. These commands have certain prerequisites, which you need to be aware of.

Before you run **definnode**, you must prepare certain information and do some manual set up.

Attention:

A node that has already been defined cannot be redefined with the **definnode** command. Including such a node in the command-line input causes the command to fail without defining any nodes. Including such a node in the node definition (*nodedef*) file, causes the definition of that node to fail, but the other nodes specified in the file are defined successfully. An error message is issued for the undefined node or nodes.

Prerequisites for running definnode or addnode

Before you run **definnode** or **addnode**, information needs to be gathered and recorded on a template similar to the example in “Appendix A. Node Attributes Template” on page 13. This information can be entered into a node definition (*nodedef* file), or it can be entered at the command line.

If you intend to use a *nodedef* file, start with the sample file in **/opt/csm/install/nodedef.sample** and fill in the information from the node-attribute planning template that you completed earlier.

See the man page or *IBM Cluster Systems Management for Linux Technical Reference* for more details about the node definition file.

Run definnode or addnode

The **definnode** command defines all the nodes in the cluster. It does not actually install the nodes. Node installation is done by **installnode**. If you run **addnode**, you do not need to run **installnode** because **addnode** runs **installnode** for you.

Note: Wherever the **definnode** command is used in the following discussion, the **addnode** command could be substituted. The **definnode** and **addnode** arguments and usage are the same.

If some arguments are not provided, the command prompts you for each piece of information that it needs. If you should inadvertently miss a required option, the command prompts you for the missing information.

You can either use a node-definition file to define the nodes, console servers, and service processors to the cluster, or you can enter the information from the command line. To use a node definition file in order to define the nodes, console server information and service processors, type:

```
definnode -f nodedef
```

To see the arguments that you need to enter from the command line, type:

```
definenode -h
```

All of the arguments are required when the command is run. The command prompts for missing information when some or all of the arguments are not provided. To use this method of input, type the command without any arguments:

```
definenode
```

See the man page or *IBM Cluster Systems Management for Linux Technical Reference* for details on **definenode** or **addnode** command-line syntax and more examples of the usage of the command. See “Example of **definenode** command run interactively” for an example that demonstrates the interactive approach.

After **definenode** has been run successfully, verify the node definitions, and then run **installnode**. See “Verifying node definitions” on page 7 for details.

Some error messages may be returned if **definenode** is not completely successful. See “FAQs, hints, tips, and troubleshooting” on page 9 for troubleshooting information.

Example of **definenode** command run interactively

If you run the **definenode** command without any options, the program prompts you for the required information. Also, if you miss a piece of required information, the program prompts you for that information.

The following example shows sample input for nodes, console servers, and service processors with an interactive program. The example uses the **definenode** command, but the **addnode** command can be used instead with the same usage and arguments. Instead of requiring you to enter everything at once on the command line, the interactive program allows you to enter a little bit at a time. User input is shown in bold type.

```
Enter starting node name (hostname or IP address): clsn01
Enter number of nodes to define (default = 1): 12
Enter list of Hardware Control Points (press ENTER for none):
  Format: hwctrlpt[:method:spname] [,...]
         mspname = Hardware Control Point hostname
                 or IP address.
         method = Power method (default=netfinity)
         spname = Starting service processor name
                 or 'hostname' (default=node01)
  Example: hwctrlpt1::node06,hwctrlpt2,hwctrlpt3
  Example: hwctrlpt1::hostname,hwctrlpt2::hostname
mgtn03,mgtn04::hostname
Enter list of Console Servers (press ENTER for none):
  Format: csname[:method:csnum:port] [, ...]
         csname = Console server name (hostname or IP address)
         method = Console method (default=esp)
         csnum = Console server number (default=1)
         port = Starting console port number (default=1)
  Example: cs1:::4,cs2:conserver,cs3
mgtn02
Enter Hardware Type (default = netfinity): netfinity
definenode: Adding CSM Nodes:
definenode: Adding Node clsn01.ppd.pok.ibm.com(9.114.133.193)
definenode: Adding Node clsn02.ppd.pok.ibm.com(9.114.133.194)
definenode: Adding Node clsn03.ppd.pok.ibm.com(9.114.133.195)
definenode: Adding Node clsn04.ppd.pok.ibm.com(9.114.133.196)
definenode: Adding Node clsn05.ppd.pok.ibm.com(9.114.133.197)
definenode: Adding Node clsn06.ppd.pok.ibm.com(9.114.133.198)
definenode: Adding Node clsn07.ppd.pok.ibm.com(9.114.133.199)
definenode: Adding Node clsn08.ppd.pok.ibm.com(9.114.133.200)
definenode: Adding Node clsn09.ppd.pok.ibm.com(9.114.133.201)
```

```
definenode: Adding Node c1sn10.ppd.pok.ibm.com(9.114.133.202)
definenode: Adding Node c1sn11.ppd.pok.ibm.com(9.114.133.203)
definenode: Adding Node c1sn12.ppd.pok.ibm.com(9.114.133.204)
```

Verifying node definitions

After **definenode** has run, the management server has been set up with all the node information for CSM. The cluster nodes are now ready to be installed. This set up, however, may not completely suit your needs. This section describes how to verify and customize the cluster definitions before the actual installing of the nodes takes place. Since the actual node installation has not happened yet, you can make changes to any node definitions here.

Verify the **csm** node information as follows:

1. Run **lsnode -P** to display whether the node is on the PreManagedNodes object list. If the node is not on the list, it will not be installed.
2. Run **lsnode -AI -P** to display all the information about each node.
3. Run **chnode -P** to change the attributes of a node if that is needed.
4. Run **rmnode -P** to remove a node before redefining it if that is needed.

If something needs to be corrected, either you can use **rmnode -P** to remove the node that was not successfully defined and then rerun **definenode** with the correct arguments, or you can use **chnode -P** to make changes to any attributes of a node. Note that all of the attributes for a node might not be filled in at this point. See the **chnode**, **definenode**, **lsnode**, and **rmnode** man pages or the *IBM Cluster Systems Management for Linux Technical Reference* for more information.

Changing internal service processor user IDs and passwords

The service processor password file is created from `/etc/opt/csm/netfinity_power.config.template` when **definenode** is run. Afterwards you can modify the `netfinity_power.config` file to specify individual passwords and user IDs for each node, if needed.

Note: Manufacturing provides a default user ID and password for the service processors. This default user ID and password is then propagated to all the nodes in the cluster when **definenode** is run. Do not modify the template file.

For more information on `netfinity_power.config`, see the man page or the *IBM Cluster Systems Management for Linux Technical Reference*.

Running installnode

This command is used to install all the nodes in the cluster by running **makenode**. The appropriate software listed in “Specified Operating Environment” on page 1 is installed automatically by the **installnode** command on the nodes if it is not already installed or if it is installed at a previous level.

The **installnode** command also displays the installation status for each node. In addition a log file is maintained in `/var/log/csm/installnode.log` that contains information on what happened during installation on each node.

For more information on **installnode** and **makenode**, see the man pages or the *IBM Cluster Systems Management for Linux Technical Reference*.

Prerequisites for running installnode

Before the **installnode** command can be run, the following must be done on the management server:

1. NFS must be available to mount and unmount `/tftpboot` remotely. The **nfs-utils** RPM contains NFS.
2. The **dsh** command must be available to perform remote commands on the nodes; that is, security must be set up on each node in such a way that **dsh** is allowed to run commands on that node. **dsh** is used to perform an NFS mount of `/tftpboot` to each node and to run **makenode** from the mounted

/ftppboot on each node. It is also used to determine the output of **makenode**, which is then recorded in the **installnode** log (**/var/log/csm/installnode.log**). The log can be reviewed to make sure that the nodes were installed successfully.

3. The **installms** and **definnode** commands must have been run successfully.
4. Verify that the node definitions are correct, or change the node definitions. See “Verifying node definitions” on page 7. When you are satisfied with the node definitions, run **installnode**.

Run installnode

1. Run **installnode**.
2. You should now have a working cluster. See “Getting started with the newly installed cluster” on page 9 to make sure that the cluster is functioning successfully.
3. See **/var/log/csm/installnode.log** for information on what happened on each node during installation if there is a problem. To determine how to handle a problem, see “FAQs, hints, tips, and troubleshooting” on page 9.

Attention:

After each node is installed by running **installnode**, you need to reboot the node to enable remote console support. The node does not have to be rebooted if remote console support is not being used.

Adding a node to an existing cluster

You can add another node to the cluster either by running **definnode** and then **installnode** again or by running **addnode** again. To add a new node to the cluster, do the following:

1. You must assign the correct host name, service processor information, and console server information. For example, you need to assign a proper unused console port number. To find out what attributes are already in use you need to see the attributes not only for managed nodes but also for PreManaged nodes because some nodes might not have been installed yet. Type the following to see the attributes of managed nodes:

```
lsnode -A1
```

And, type the following to see the attributes of PreManaged nodes:

```
lsnode -A1P
```

2. To define the node, type:

```
definnode
```

You will be prompted for the required information.

3. After **definnode** is completed successfully, run **installnode**.
4. To verify that the node is installed, type:

```
lsnode -A1
```

Removing a node from an existing cluster

Removing a node from a cluster does not uninstall CSM and its prerequisites from the node. Rather, it disassociates the node from its management server. It removes the node from the database of the management server, and it informs the node that it is no longer attached to the management server. To remove a node from the cluster, type:

```
rmnode hostname
```

A removed node can be added back into the cluster by running **definnode** or **addnode** again.

Getting started with the newly installed cluster

This section tells you how you can determine whether the installation was successful. It also gives you some suggestions on how to get started using Cluster Systems Management. After installation is successfully completed, remote RMC and CSM commands are enabled. To verify that the installation was successful, enter the following commands. If everything is as it should be, you should see the following results:

- If you enter the CSM command **lsnode** on the command line, a list of the active nodes in the cluster is displayed.
- To verify that **dsh** is working on all of the nodes by running the **date** option, type:

```
dsh -a date
```

A list of nodes with the date on each node is returned.

- To verify the power status of the nodes (whether they are on or off), type:

```
rpower -a query
```

A list of nodes with their associated power state is returned.

- To verify whether the nodes are reachable, type:

```
lsnode -p
```

The ping status of the nodes is returned.

- To receive a wall message on the management server when a node is shut down, type:

To try out monitoring, use the following example.

- To set up a condition that causes a wall message to be broadcast on the management server when a node is shut down, type:

```
startcondresp NodeReachability BroadcastEventsAnyTime
```
- To show the conditions that are currently monitored, type:

```
lscondition
```
- To show the response associated with the condition, type:

```
lscondresp
```

FAQs, hints, tips, and troubleshooting

This section has frequently asked questions that can help to troubleshoot problems or give hints and tips on how to do something more easily or efficiently. The first group of questions are general troubleshooting questions. They are followed by a special group of questions on how to handle the RMC ACL file. In addition, for problems with the monitoring function, see the *Diagnostics* chapter in *IBM Cluster Systems Management for Linux Monitoring HOWTO*.

Question What troubleshooting tools are available to me?

Answer You can always specify the verbose option (**-v**) on any of the installation commands and rerun the command to get more information on what the problem is.

Question What do I do when a node that I want to specify for the cluster is said to be already defined?

Answer Remove the node from the cluster (by using **rmnode** or **rmnode -P**) and then define it again (by using **definenode**) , or use **chnode -P** to change the attributes of the node.

Question What do I do if **installms** does not complete successfully?

Answer

Rerun the command and specify the verbose option (**-v**). This will show you where the problems occurred. The verbose option is available for this purpose on all the installation commands.

Question

What do I do if **installnode** does not complete successfully?

Answer

Consult the **/var/log/csm/installnode.log** log. Everything that occurred on each node during installation is recorded in this log.

Question

What do I do if I have one or more nodes that are displayed as ManagedNodes by the **lsmode** command and as PreManagedNodes by the **lsmode -P** command?

Answer

Ensure that the **ctrmc** daemon is running on the management server by running the following:

```
lssrc -s ctrmc
```

If the output shows that **ctrmc** is inoperative, run the following:

```
startsrc -s ctrmc
```

Then verify that **ctrmc** is now active by running the following again:

```
lssrc -s ctrmc
```

After **ctrmc** is active, you can run **installnode** again on the management server. This moves the PreManaged Nodes completely to ManagedNodes and finishes the necessary processing.

Question

How do I update the cluster if I change the host name of the management server?

Answer

From the management server, use **dsh** to run the **mgmtsvr** command. For example:

```
dsh -a mgmtsvr
```

Note: You may have to update the **dsh** access files (for example, the **.rhosts** file) because the management server host name has changed.

Question

How do I update the cluster if I change the host name of a managed node?

Answer

After you change the host name on the managed node, run the **chnode** command on the management server for this managed node to set the new host name.

Question

I am running into problems with the CSM installation process, and I am getting a message that the device is busy. What do I do?

Answer

Check that you are running the **installms** command from a temporary directory and not from the CD-ROM. Running **installms** from the CD-ROM causes a device-busy condition because it interferes with the mounting and unmounting of the Red Hat CD-ROMs.

ACL File FAQs

The RMC ACL file is located in **/var/ct/cfg/ctrmc.acls**. The management server uses the RMC ACL file as its authorization mechanism. You may want to update the RMC ACL file to manually fix problems during installation. The following questions and answers can guide you through the process of updating the ACL file to provide access to the resource classes on the management server machine to managed nodes.

Note: Any time that the ACL file is updated, the **refresh** command needs to be issued in order for the updates to take effect as follows:

```
refresh -s ctrmc
```

Question If the RMC ACL file does not exist and I need to add entries to the ACL file, what do I do?

Answer Copy the sample **ctrmc.acls** file from **/usr/sbin/rsct/cfg** to the **/var/ct/cfg** directory. Study the file to learn how to add entries for specific resource classes. To make sure that the new ACL file takes effect, type:

```
refresh -s ctrmc
```

Question How do I modify an ACL file to provide read/write access to the root user and read access to all other users on the remote node **hostname1** to the following resource classes:

- IBM.PreManagedNode
- IBM.ManagedNode
- IBM.NodeGroup

Answer Open the existing **ctrmc.acls** file using a text editor, and add the following stanzas:

```
IBM.PreManagedNode
  root@hostname1 * rw
  hostname1      * r

IBM.ManagedNode
  root@hostname1 * rw
  hostname1      * r

IBM.NodeGroup
  root@hostname1 * rw
  hostname1      * r
```

Save and close the file. Issue the **refresh** command:

```
refresh -s ctrmc
```

The output when the **ctrmc.acls** file is listed should show the following:

```
IBM.PreManagedNode
  root@hostname1 * rw
  hostname1      * r

IBM.ManagedNode
  root@hostname1 * rw
  hostname1      * r

IBM.NodeGroup
  root@hostname1 * rw
  hostname1      * r

OTHER
  root@LOCAHOST * rw
  LOCALHOST     * r
```

Question How do I modify the **ctrmc.acls** file to provide read/write access to the root user and read access to all other users on a remote node **hostname2** to the following resource classes:

- IBM.PreManagedNode
- IBM.ManagedNode
- IBM.NodeGroup

Answer Open the existing **ctrmc.acls** file using a text editor, and add the following entries to the existing stanzas:

```
IBM.PreManagedNode
  root@hostname2 * rw
  hostname2      * r
```

```
IBM.ManagedNode
  root@hostname2 * rw
  hostname2     * r
```

```
IBM.NodeGroup
  root@hostname2 * rw
  hostname2     * r
```

Save and close the file. Save and close the file. Issue the **refresh** command:

```
refresh -s ctrmc
```

The output when the **ctrmc.acls** file is listed should show the following:

```
IBM.PreManagedNode
  root@hostname1 * rw
  hostname1     * r
```

```
  root@hostname2 * rw
  hostname2     * r
```

```
IBM.ManagedNode
  root@hostname1 * rw
  hostname1     * r
```

```
  root@hostname2 * rw
  hostname2     * r
```

```
IBM.NodeGroup
  root@hostname1 * rw
  hostname1     * r
```

```
  root@hostname2 * rw
  hostname2     * r
```

```
OTHER
  root@LOCAHOST * rw
  LOCALHOST    * r
```

Appendix B. Notices

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

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

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

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

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 other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

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
Department LJEB/P905
2455 South Road
Poughkeepsie, NY 12601-5400
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.

COPYRIGHT LICENSE:

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

Trademarks

The following trademarks apply to this book:

IBM and AIX are registered trademarks of International Business Machines Corporation.

Linux is a registered trademark of Linus Torvalds.

Red Hat and RPM are trademarks of Red Hat, Inc.

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

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

Publicly Available Software

IBM Cluster Systems Management for Linux includes software that is publicly available:

cfengine	A software package that is licensed under GPL and is used to create customization scripts.
Conserver	An application that adds logging and multi-user access for remote administration of serial ports, using locally installed multi-port serial interfaces and/or "reverse-telnet" to console servers.
DBD-CSV, DBI	Licensed under GPL or Artistic, these are dynamically loaded Perl modules.

fping	Licensed under BSD, this is executed as a separate binary.
Perl	Practical Extraction and Report Language is licensed under the Artistic license.
Pidentd	Public domain program by Peter Eriksson that implements the RFC-1413 identification server.
SQL-Statement	Licensed under GPL or Artistic, this is a dynamically loaded Perl module.

This book discusses the use of these products only as they apply specifically to the IBM Cluster Systems Management for Linux product.

Note: The distribution for these products includes the source code and associated documentation. All copyright notices in the source code and the documentation must be respected. You can find version and distribution information for each of these products that are part of your selected install options in the README file.

Index

Special Characters

\$MANPATH variable 2
\$PATH variable 2
/ftpboot partition 3

A

about this book v
ACL file FAQs 10
adding a node to an existing cluster 8
addnode command 5
 example 6
 prerequisites for 5
audience of this book v

C

changing user IDs and passwords
 for internal service processors 7
configuring the remote shell 3
creating the /ftpboot partition 3

D

definenode command 5
 example 6
 prerequisites for 5
disk space requirements 2
dsh
 preparing for 3
 setting up security for 3

E

example of definenode or addnode run interactively 6

F

FAQs 9
 for the ACL file 10

G

getting started with the newly installed cluster 9

H

hardware requirements 1
 for remote control 1

I

identd service 4
installation processes 4
installing all the nodes in the cluster 7
installing the management server 4
installms command
 prerequisites for 4

installnode command 7
 prerequisites for 7
installnode log
 used for debugging 7

L

limitations 2

M

management server
 installing 4
memory requirements 2

N

netfinity_power.config file 7
network requirements 2
node attribute template 13
node definitions
 verifying 7
nodedef file 5

O

obtaining publications vi
overview of installation processes 4
overview of set-up process 1

P

planning for CSM 2
 prerequisites for running addnode or definenode 5
 prerequisites for running installnode 7
pre-installation tasks
 configuring the remote shell 3
 preparing for dsh 3
 preparing node-attributes table 2
 preparing nodedef file 5
prerequisite knowledge for this book v
prerequisites
 for running addnode 5
 for running definenode 5
 for running installms 4
 for running installnode 7

R

remote console support 4
remote control
 hardware requirements 1
 network considerations 2
 planning for 2
removing a node from an existing cluster 8
requirement
 for rebooting to enable remote console support 4
 to use fully qualified host name for dsh 3

- requirements 1
- running addnode 5
 - prerequisites for 5
- running definenode 5
 - prerequisites for 5
- running installms
 - prerequisites for 4
- running installnode 7
 - prerequisites for 7

S

- setting the \$MANPATH variable 2
- setting the \$PATH variable 2
- software requirements 1
- specified operating environment 1

T

- troubleshooting information 9

V

- verifying node definitions 7

W

- web site for CSM vi

Readers' Comments — We'd Like to Hear from You

IBM Cluster Systems Management for Linux®
Set-Up HOWTO
Version 1 Release 1

Publication No. SA22-7853-00

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	<input type="checkbox"/>				

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	<input type="checkbox"/>				
Complete	<input type="checkbox"/>				
Easy to find	<input type="checkbox"/>				
Easy to understand	<input type="checkbox"/>				
Well organized	<input type="checkbox"/>				
Applicable to your tasks	<input type="checkbox"/>				

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? Yes No

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

Name

Address

Company or Organization

Phone No.

Readers' Comments — We'd Like to Hear from You
SA22-7853-00



Cut or Fold
Along Line

Fold and Tape

Please do not staple

Fold and Tape

PLACE
POSTAGE
STAMP
HERE

IBM Corporation
Department 55JA, Mail Station P384
2455 South Road
Poughkeepsie NY 12601-5400

Fold and Tape

Please do not staple

Fold and Tape

SA22-7853-00

Cut or Fold
Along Line



Program Number: 5799-GNJ

SA22-7853-00

