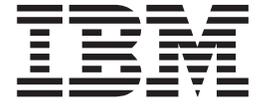


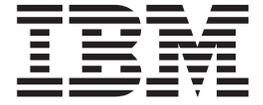
IBM Cluster Systems Management for Linux[®]



Overview HOWTO

Version 1 Release 1

IBM Cluster Systems Management for Linux[®]



Overview HOWTO

Version 1 Release 1

Note!

Before using this information and the product it supports, read the information in "Notices" on page 7.

First Edition (June 2001)

This edition of the *IBM Cluster Systems Management for Linux Overview 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: mhvrcfs@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	Monitoring and Controlling Hardware	3
Who Should Use This HOWTO	v	Running Commands Remotely	3
Typographic Conventions	v	Configuration File Manager	3
Related Information	v	Monitoring System Events	4
How to Obtain Publications	vi	Security	4
		Diagnosing Problems	5
IBM Cluster Systems Management for Linux Overview	1	Notices	7
Setting Up IBM Cluster Systems Management for Linux	1	Trademarks	8
Managing Node and Node-Group Information	2	Publicly Available Software	9

About This HOWTO

This HOWTO is an overview of the the IBM Cluster Systems Management for Linux set of tools. It briefly describes the administrative tasks that can be accomplished with greater ease and efficiency by using this set of tools and then directs you to detailed related information for each task.

Who Should Use This HOWTO

This HOWTO is intended for system administrators who want to use Cluster Systems Management for Linux. 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 Remote Control HOWTO, SA22-7856-00

IBM Cluster Systems Management for Linux Set-Up HOWTO, SA22-7853-00

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.

A README 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 Overview HOWTO*, csmovrvw.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 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>

IBM Cluster Systems Management for Linux Overview

IBM Cluster Systems Management for Linux (CSM) provides a distributed system management solution for machines, or *nodes*, that are running the Linux operating system. With this software, an administrator can easily set up and maintain a Linux cluster by using functions like monitoring, hardware control, and configuration file management. The concepts and software are derived from IBM Parallel System Support Programs for AIX (PSSP) and from applications available as open source tools.

Specifically, within the cluster, nodes can be added, removed, changed, or listed (with persistent configuration information displayed about each node in the list). Commands can be run across nodes or node groups in the cluster, and responses can be gathered. Nodes and applications can be monitored as to whether they are up or down; CPU, memory, and system utilization can be monitored; and automated responses can be run when events occur in the cluster. Configuration File Manager is provided for synchronization of files across multiple nodes. A single management server is the control point for the CSM cluster.

Note that CSM manages a loose cluster of machines. It does not provide high availability services or fail-over technology although high-availability clusters can be part of the set of machines that CSM is managing.

More information is provided on these tasks as follows:

1. Setting up IBM Cluster Systems Management for Linux
2. Managing node and node group information
3. Monitoring and controlling hardware
4. Running remote commands on multiple nodes
5. Monitoring system events
6. Managing and synchronizing configuration files for all nodes
7. Providing security
8. Diagnosing problems

Setting Up IBM Cluster Systems Management for Linux

The *IBM Cluster Systems Management for Linux Set-Up HOWTO* provides a simple process for installing and configuring CSM on an existing Linux system. This process allows you to do the following:

1. Install IBM RSCT and CSM code on the management server.
2. Automatically add nodes to the cluster during the installation process.
3. Install and configure IBM RSCT and IBM Cluster Systems Management for Linux on all nodes in a cluster from a single management server.

For more information, see the man pages or *IBM Cluster Systems Management for Linux Technical Reference* for the following set up commands and files:

1. **installms**
Installs the management server
2. **addnode**
Can be used instead of **definnode** and **installnode** for suitable installations

3. **definenode**
Gathers all the information necessary to install the nodes
4. **installnode**
Installs the nodes and brings up the necessary servers on them
5. *nodedef* file
Node definition file for cluster nodes

Managing Node and Node-Group Information

The distributed management server provides a set of commands for managing nodes. It stores information about nodes in a central repository, and it defines static and dynamic node groups. These definitions are then accessible to the Configuration File Manager (**cfm**) command for configuration file management, the **dsh** command for running shell commands remotely, for hardware control, and for monitoring the cluster by using the Event Response subsystem (ERRM). All of these functions rely on the definitions stored by the **nodegrp** command. Thus, a node group is defined in only one place and is then accessible for use by other functions.

Persistent information on each node is kept, including operating system type, host name, machine type, model, and serial number. In addition, the status of the node is determined periodically by means of the **fping** command.

The node and node group commands are built on top of a Perl DBI layer backed by a set of DBDs (database drivers) so that data can be stored in a variety of formats and shared with other tools.

See the man pages or the *IBM Cluster Systems Management for Linux Technical Reference* for details on the following commands that manage node and node-group information:

1. **addnode**
Adds a node to the CSM cluster database.
2. **chnode**
Changes an attribute of a node in the CSM cluster database.
3. **dmsctrl**
Allows **fping** and power status parameters to be changed.
4. **lsnode**
Displays information about the nodes in the CSM cluster, for example, the cached status on whether the node has been reachable.
5. **nodegrp**
Defines node groups within the CSM cluster for use by other functions such as the configuration file manager, the **dsh** command, the event response subsystem, and the hardware control commands.
6. **rmnode**
Removes a node from the CSM cluster database.
7. **whichdb**
Specifies the repository for node information.

Monitoring and Controlling Hardware

You can control the hardware on remote nodes by using the remote control commands. For example, you can control computers on a ship from an office on the mainland, provided the correct connectivity exists.

See *IBM Cluster Systems Management for Linux Remote Control HOWTO* for details on how to set up remote power control. See the man pages or *IBM Cluster Systems Management for Linux Technical Reference* for details on the following commands:

1. **rconsole**

Opens a remote console.

2. **rpower**

Boots and resets hardware, powers hardware on and off, and queries the power state; for example, the **resetsp** option resets the service processor.

Running Commands Remotely

The distributed shell (**dsh**) command runs commands remotely across multiple nodes. It optionally can use any underlying remote shell that is specified by the user (for instance, a remote secure shell that complies with the IETF (Internet Engineering Task Force) Secure Shell protocol). By default, **rsh** is used.

Note: It is the system administrator's responsibility to configure and enable remote shell access to other systems.

The **dsh** command can retrieve a complete list of the nodes in the CSM cluster or the list of nodes in a specified node group.

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.

See the man pages or *IBM Cluster Systems Management for Linux Technical Reference* for details on the following commands:

1. **dsh**

Issues remote shell commands and the options associated with them to multiple nodes.

2. **dshbak**

Presents formatted output from the **dsh** command.

Configuration File Manager

Configuration File Manager provides a file repository for the common configuration files among nodes in a cluster. In general, all the configuration files that need to be shared are stored in one location on the management server. Changes to these files are propagated and synchronized throughout the cluster. Though the files are common, there are mechanisms to allow for variations based on groups, IP address, and host name.

Configuration File Manager is built on top of the GNU software package cfengine. The cfengine software package is a scripting package that uses a class-based decision structure to test and configure UNIX-like systems attached to a TCP/IP

network. There are many capabilities built into cfengine itself, which a system administrator can use over and above what Configuration File Manager uses.

Configuration File Manager greatly enhances the copy functionality and usability of cfengine by providing the concept of a repository. Instead of requiring an administrator to write a cfengine script to keep files up to date, the repository allows automatic updating without script changes.

Note: CFM is based on cfengine which currently uses a host list file (`/etc/opt/csm/cfd.conf`) to determine which hosts can access the configuration files that it controls. This may be a security vulnerability due to IP-address spoofing or a compromised DNS. This possible security vulnerability will be addressed in the next release.

See the man pages or *IBM Cluster Systems Management for Linux Technical Reference* for more details on the **cfm** and **cfonce** commands.

At the time this document was written, detailed information on cfengine could be found at the following URL: <http://www.iu.hioslo.no/cfengine>

Monitoring System Events

A flexible distributed system monitoring application is provided by CSM. This monitoring application allows the administrator to define conditions of interest to monitor on the system. An event occurs when a monitored condition of interest reaches a threshold that is defined in an event expression. When an event occurs, automated responses to that event take place. Many actions can be defined as part of these responses, including notification, running a predefined response script, or running a user-defined script.

A full set of commands is provided to tailor this application to your needs. In addition, predefined conditions and responses are provided for easy implementation so that you can get up and running quickly and easily. This rich set of predefined conditions and responses can be used directly or can be taken as examples to be copied and modified. Among the system resources that can be monitored are:

- File systems
- Programs
- System resources
- Node availability
- Other resources by means of sensors

The application, its components, and the predefined conditions and responses are fully described in the *IBM Cluster Systems Management for Linux Monitoring HOWTO*. The commands are available as man pages and are also compiled for easy reference in the *IBM Cluster Systems Management for Linux Technical Reference*.

Security

Security on a single system is provided by the operating system in that only root can run or modify functions. Flexibility is provided for the degree of security required by the specific environment because remote shells that conform to the IETF (Internet Engineering Task Force) Secure Shell protocol can be specified by using the **dsh** command for the appropriate situations. Network security for other functions is built on the **identd** function.

See the *IBM Cluster Systems Management for Linux Monitoring HOWTO* for details on authorization and the **dsh** man page or the *IBM Cluster Systems Management for Linux Technical Reference* for details on how to specify the remote shell of your choice by using the DSH_REMOTE_CMD environment variable.

Diagnosing Problems

Cluster Systems Management (CSM) makes use of several other tools. It is helpful to understand the relationship between CSM and these tools in order to diagnose problems. The tools that CSM uses are described in the following table:

Tool	What It Does
Perl DBI package	Stores database information in a variety of formats
Resource Monitoring and Control (RMC) subsystem	Monitors conditions and communicates with all nodes. RMC needs to be running on each node, and the security access control list (ACL) file needs to allow the nodes to communicate with the management server. See "Security Considerations" in the "Overview" chapter of the <i>IBM Cluster Systems Management for Linux Monitoring HOWTO</i> .
dsh	Runs some commands on the nodes. Security needs to be set up on each node to allow this for the remote shell that is used by dsh . The default remote shell is rsh .
fping	Periodically gets the status of each node
cfengine	Transfers files for the Configuration File Manager

Here are a few tips to help diagnose problems with a CSM cluster:

- To ensure that the database attributes are correct for each node, type:
`lnode -A1`
- To list the status of the RMC daemons, type:
`lssrc -a`
- To review the audit log for monitoring events, type:
`lsaudrec`
- If you have modified the RMC access control list (ACL) file, make sure that it is correct on each node. The RMC ACL file, when it exists because the default permissions have been modified, is located at `/var/ct/cfg/ctrmc.acls`.
- Make sure the cfengine security file, `/etc/opt/csm/cfd.conf`, on the management server contains all the nodes.
- Examine the cfengine log file.
- If you are using the **rsh** remote shell, make sure that the `/root/.rhosts` file on each node contains the management server.

See the "Diagnosis Information" chapter and the "Security Considerations" section of the "Overview" chapter in the *IBM Cluster Systems Management for Linux Monitoring HOWTO* for troubleshooting hints and tips and for detailed information on authorization and the ACL file respectively. See the ACL File FAQ in the *IBM Cluster Systems Management for Linux Set-Up HOWTO* for information on troubleshooting the RMC ACL file.

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 under the GPL license 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.

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

IBM Cluster Systems Management for Linux®
Overview HOWTO
Version 1 Release 1

Publication No. SA22-7857-00

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
SA22-7857-00



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

SA22-7857-00

Cut or Fold
Along Line



Program Number: 5799-GNJ

SA22-7857-00

