IBM Cluster Systems Management for Linux®

Planning and Installation Guide

Version 1.1

SA22-7853-02

IBM Cluster Systems Management for Linux®

Planning and Installation Guide

Version 1.1

SA22-7853-02

Note!
Before using this information and the product it supports, read the information in “Notices” on page 37.

Third Edition (December 2001)

This edition of the *IBM Cluster Systems Management for Linux Planning and Installation Guide* applies IBM Cluster Systems Management for Linux Version 1.1, program number 5765–E88, 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.

Contents

About This Book

This book is the updated version of the *IBM Cluster Systems Management for Linux Set-Up HOWTO* (SA22-7853-00) that was last published in June, 2001. This book provides planning, installation, backup and recovery, and troubleshooting information related to getting the Cluster Systems Management product up and running. The information includes the specified operating environment for the IBM Cluster Systems Management for Linux (CSM) suite of tools and describes how to install and set up a CSM cluster in two environments, either on an existing group of nodes or as a full installation of both the Linux operating system and CSM on the nodes of the cluster.

Who Should Use This Book

This guide 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.

How to Use This Book

This book contains information to help you install IBM Cluster Systems Management for Linux. It should be read in its entirety before you begin the installation process to ensure that everything will go smoothly and easily.

Highlighting

The following highlighting conventions are used in this book:

Bold	Identifies commands, subroutines, keywords, files, structures, directories, and other items whose names are predefined by the system. Also identifies graphical objects such as buttons, labels, and icons that the user selects.
Italic	Identifies parameters whose actual names or values are to be supplied by the user.
`monospace`	Identifies examples of specific data values, examples of text similar to what you might see displayed, examples of portions of program code similar to what you might write as a programmer, messages from the system, or information you should actually type.

ISO 9000

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

Related Information

The following references contain more information about IBM Cluster Systems Management for Linux:

- *IBM Cluster Systems Management for Linux Administration Guide*, SA22–7873–00
- *IBM Cluster Systems Management for Linux Remote Control Guide and Reference*, SA22–7856-01
- *IBM Cluster Systems Management for Linux Technical Reference*, SA22–7851–01

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 Administration Guide*, am7LXadm.pdf
- *IBM Cluster Systems Management for Linux Planning and Installation Guide*, am7LXstp.pdf
- *IBM Cluster Systems Management for Linux Remote Control Guide and Reference*, am7LXrem.pdf
- *IBM Cluster Systems Management for Linux Technical Reference*, am7LXcmd.pdf

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

`http://www.ibm.com/eserver/clusters/linux`

How to Reach Us by E-mail

If you would like to contact us by e-mail, send your comments to cluster@us.ibm.com.

Chapter 1. Setting Up IBM Cluster Systems Management for Linux

The planning and installation process described in this guide helps a system administrator to 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 either on existing Linux systems or by doing a full installation of the Linux operating system and CSM on the nodes of the cluster (not the management server, which must be set up first).

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 either on an existing group of nodes (CSM-only installation) or for a full installation of both the operating system and CSM (Kickstart installation). 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 minimum hardware and software that are required for the Cluster Systems Management specified operating environment. For more detailed information, see the announcement.

Hardware Requirements

This product is supported on the IBM @server Cluster 1300 Model 7080 Primary Rack, 7080 Expansion Rack, 7081, and 7082.

Memory and Disk Space

For the management server, a minimum of 128MB of memory and 120MB of disk space is required for installing CSM, and an additional 1.5GB is required for a full installation.

For the managed node, a minimum of 128MB of memory and 20MB of disk space is required for CSM and the appropriate amount of additional disk space for the operating system and RPMS that you choose to install.

Network Requirements

TCP/IP and an Ethernet 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.

Remote Control Hardware Requirements

To support remote control, the following hardware is required:

- One IBM Netfinity® Remote Supervisor PCI adapter for every ten x330 nodes; one IBM Remote Supervisor Adapter (RSA) for every x342 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 Guide and Reference*.

Limitations

A cluster of up to 128 nodes is supported.

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 non-IBM 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.

CSM File Sets

The following file sets comprise the IBM Cluster Systems Management for Linux product:

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

Required non-IBM Software

The Red Hat Linux 7.1 distribution is required. The Red Hat 7.1 distribution includes the following required non-IBM software:

1. Expect 5.31-53 (on management server only)
2. dhcp 2.0 (on management server only)
3. glibc 2.2
4. libstdc++ 2.9
5. pdksh 5.0
6. Pidentd 3.0
7. perl 5.00503
8. make 3.7
9. nfs-utils 0.1 (on management server only)

In addition, IBM provides the following required non-IBM software:

1. atftp 0.3
2. cfengine 1.6.2
3. DBD-CSV 0.1024
4. DBI 1.14
5. fping 2.2.b.1
6. SQL-Statement 0.1016
7. SYSlinux 1.53
8. Text-CSV_XS 0.21

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

Compatibility with General Parallel File System

It is mandatory that the CSM client RPM be installed either on all of the General Parallel File System for Linux (GPFS) nodes or on none of them. The Resource Monitoring and Control (RMC) subsystem requires that the definitions of its resource classes be the same on all nodes within the GPFS cluster. If the CSM client RPM is installed on only some of the GPFS nodes, some of the RMC daemons will not be able to join their peer group successfully, an undesirable situation. The CSM management server node should not be part of the GPFS cluster because it adds other resource classes to RMC, and it is not practical or desirable to install the CSM server RPM on all of the GPFS nodes.

GPFS and CSM should be using the same level of RSCT code. The current release of RSCT is Version 2.2.

Planning for Cluster Systems Management

This section describes the tasks that you need to perform before installing this product. Before you do anything else, ensure that the required IBM software is installed.

Read *IBM Cluster Systems Management for Linux Remote Control Guide and Reference* before beginning the installation process. *IBM Cluster Systems Management for Linux Remote Control Guide and Reference* contains illustrated information on hardware and networking set up and configuration, including the default installation values; for example, **eth0**, which is the value for the first Ethernet adapter that is connected to the cluster VLAN. It also provides a sample completed node-attribute table that you can use for guidance in completing your own blank template. A blank template is provided in “Appendix A. Node Attributes Template” on page 25 for your convenience.

Pre-installation tasks

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

1. For authentication and authorization of RMC requests, CSM software components require that any reference to an IP address or short hostname of any CSM management server or managed node be consistently resolvable by DNS or **/etc/hosts** to its canonical fully-qualified domain name (FQDN), which is normally the long version of the hostname. Two specific cases are considered:
   - Where **/etc/hosts** is used before DNS, **/etc/hosts** must be set up so the hostname returned for any CSM management server or managed node is its canonical FQDN.
   - Where there are multiple alias FQDNs for the same IP address, **/etc/hosts** must be used before DNS and must be set up so the hostname returned for any CSM management server or managed node is its canonical FQDN. This set up handles situations where DNS’s sometimes return an alias instead of the canonical FQDN.
2. Some system administrators may choose to make **/usr/local** available to CSM cluster nodes as a read-only NFS mounted file system. When CSM is installed on the management server and the nodes, **cfengine** gets installed in **/usr/local**. If **/usr/local** is NFS mounted read-only, the installation of **cfengine** will be unsuccessful. If this is the case, choose one of the following two methods to avoid this problem:
   - Use the CSM management server as the NFS server for **/usr/local** for the nodes. This will cause **cfengine** to get installed in the correct place in **/usr/local** when the management server is installed. Then when the nodes are installed, it will recognize that **cfengine** is already in **/usr/local** and not try to install it again.
   - If **/usr/local** is served by another file server, follow these instructions before running **installms** and **installnode**:
     a. Copy the **cfengine** tarball from the ″**reqs/**″ subdirectory of the CSM CD to a temporary directory on the file server. (If CSM has previously been installed on the management server, the **cfengine** tarball is also in **/tftpboot/tarball**.)
     b. Gunzip and untar the tarball.

c. In the **cfengine** directory, run:

```
configure
make
make install
```

3. Information needs to be gathered to complete the template in “Appendix A. Node Attributes Template” on page 25. See the section ″Hardware Configuration ″ in *IBM Cluster Systems Management for Linux Remote Control Guide and Reference* for a detailed discussion of hardware configuration and set-up considerations and for an example of a completed node-attribute table.
4. 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.
5. IBM suggests that you create a separate partition called **/tftpboot** on the management server to contain a copy of the packages on the installation media before you install CSM. See “Creating the /tftpboot partition” for the procedure that shows how to do this.
6. IBM suggests that you add **/opt/bin** to your $PATH variable and **/opt/man** to your $MANPATH variable on the management server and on the managed nodes.
7. Prepare for the use of the **dsh** command. See “Preparing for dsh and configuring the remote shell” on page 5.
8. For a full installation, set the boot order of the nodes to boot over the network. The recommended boot order is:
   a. Diskette
   b. CD-ROM
   c. Network
   d. Hard disk

   During the full installation process, Kickstart installs the operating system on each node. The Kickstart post-installation script does any additional set-up. As the node boots for the first time, the **makenode** command is run, which converts the node from the PreManagedNode class to the ManagedNode class.

Creating the /tftpboot partition

Before installing CSM, create a partition called **/tftpboot** that consists of 1GB of space for a full installation, which includes the operating system, or 100MB for a CSM-only installation. This will hold the required RPM and tarball packages for installation.

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

1. Log in as root.
2. Run **cfdisk.**
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) for a CSM-only installation. Enter `1500MB` for a full installation.
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  /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 **csm.dsh** 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**. If **rsh** is used, it is automatically configured on the nodes during a full installation. (Note that **ssh** and Kerberos are not automatically configured on the nodes.)
2. The following steps are necessary for a CSM-only installation only.
   a. 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.
   b. 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.
   c. 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. The full install process configures **rsh** on the nodes if DSH_REMOTE_CMD is set to **rsh** (or if DSH_REMOTE_CMD is not set). This configuration adds the management server to **/root/.rhosts**, starting the **rsh** daemon on the nodes.
   d. 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 set-up 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*.

Note: The Distributed Command Execution Environment (DCEM) graphical user interface is provided to make it easy to use the **dsh** command functions. See the **dcem** man page and *IBM Cluster Systems Management for Linux Administration Guide* for information on how to start and how to use DCEM.

Overview of Cluster Systems Management installation

Two installation processes are offered to install the nodes. If the Linux operating system is not installed, you can install both the Linux operating system and Cluster Systems Management on the nodes. This is referred to as a full installation. If a Linux operating system already exists, you can install just Cluster Systems Management. This is referred to as a CSM-only installation.

For both a full installation and a CSM-only installation, 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** to install the management server.
2. Run **addnode** to install the nodes.

   The **addnode** command runs **definenode**, **setupks** (if this is a full installation), and then **installnode** automatically.
3. Run **monitorinstall** to view the results and status of the installation.

If you need more control and would like the ability to double-check and make interim changes during installation, run the underlying commands individually as follows:

For a CSM-only installation, do the following:

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.
4. Run **monitorinstall** to view the results and status of the installation.

For a full installation, do the following:

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 **setupks** command. You can use the **setupks** command to tailor your configuration. See the **setupks** man page and the annotated **kscfg.tmpl** file in the Appendix for more information.
4. Run the **installnode** command.
5. Run the **monitorinstall** command to view the results and status of the installation.

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:

For a CSM-only installation, 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 full installation process does a reboot automatically.

Running installms

The **installms** command performs the tasks that are necessary to make this system a management server. It installs the Open Source and IBM CSM 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 CSM, RPM, 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 back level) and then installs or replaces the required packages, as necessary.

After **installms** has been run successfully, the management server is installed. You are ready to run the **definenode** or **addnode** command next.

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

Running definenode or addnode

After the management server is installed by running **installms**, run **definenode** 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 **definenode**, you must prepare certain information and do some manual set up.

A node that has already been defined cannot be redefined with the **definenode** 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 definenode or addnode

Before you run **definenode** or **addnode**, information needs to be gathered and recorded on a template similar to that in “Appendix A. Node Attributes Template” on page 25. A completed example of this table can be found in *IBM Cluster Systems Management for Linux Remote Control Guide and Reference*. This information can be entered into a node definition (**nodedef**) file, or it can be entered at the command line.

A **nodedef** file allows you enter the host names of the nodes that you wish to define in a file and thus avoid the error-prone task of having to type all of the host names on the command line in order to define nodes. If you intend to use a **nodedef** file, start with the sample file in **/opt/csm/install/nodedef.sample** and complete the information from the node-attribute planning template that you completed earlier.

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

Attention:
IBM suggests using the **nodedef** file and not the command line to define the nodes in the cluster.

Run definenode or addnode

The **definenode** 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 **definenode** command is used in the following discussion, the **addnode** command could be substituted. Unless otherwise noted, the **definenode** and **addnode** arguments and usage are similar.

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:

```
definenode -f nodedef
```

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

```
definenode -h
```

To define the node with the host name **clsn02.ppd.pok.ibm.com** to the cluster, type:

```
definenode -n clsn02.ppd.pok.ibm.com
```

The command prompts for missing information when some or all of the arguments are not provided.

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 9 for details.

Some error messages may be returned if **definenode** is not completely successful. See “Chapter 4. FAQs, hints, tips, and troubleshooting” on page 21 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][,...]
                 hwctrlpt = 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
Enter Install Method (csmonly or kickstart, default=csmonly):  kickstart
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 clsn10.ppd.pok.ibm.com(9.114.133.202)
definenode: Adding Node clsn11.ppd.pok.ibm.com(9.114.133.203)
definenode: Adding Node clsn12.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.

For a CSM-only installation, the cluster nodes are now ready to be installed.

For a full installation, **setupks** must still be run before verification can take place. Perform the activities in “Running setupks” on page 10 before attempting to install the cluster nodes for a full installation.

When you are ready, this section describes how to verify and customize the cluster node 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 -Al -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 completed 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.

ISP password file generated when definenode is run

The **definenode** command is run to set up the hardware-control-point and console-server attributes that are needed by the **rpower** and **rconsole** commands. Before running **setupks** to perform a full installation, you need to make sure the internal service processor (ISP) passwords are correct.

The ISP password file is created from **/etc/opt/csm/netfinity_power.config.templ** 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. Or, you can leave the default passwords if that is suitable.For more information on **netfinity_power.config**, see the man page or the *IBM Cluster Systems Management for Linux Technical Reference*.

More information and additional steps on changing the user IDs and passwords are supplied in *IBM Cluster Systems Management for Linux Remote Control Guide and Reference*.

Running setupks

The **setupks** command collects configuration information and uses a template create the Kickstart configuration file for each node containing the information that it has collected. The command also:

- copies Red Hat CD-ROMs
- creates a **/etc/dhcpd.conf** file
- creates a **firstboot** script for each node
- creates a **prelinux** configuration file for each node
- collects MAC addresses for all nodes.

A sample annotated template is located in the Appendix. The current template is located on your system at **/opt/csm/install/kscfg.tmpl**.

You can use the template as-is or modify it. See the annotated **kscfg.tmpl** file in the Appendix for instructions on how to modify this file. Modifications made to the template file affect the entire cluster and should be made before running **setupks**.

You can also modify the generated Kickstart configuration file for each node. The resulting Kickstart configuration file that is generated for each node by **setupks** is called **/tftpboot/ks71/***node-ip-address***-kickstart** Modifying this file affects only the settings on this node and should be done after running **setupks**.

In particular, there are variables in the Kickstart configuration template and the generated file in the format #VARIABLE# that must not be deleted. Many of these are automatically customized with the appropriate values during the process of generating the Kickstart configuration file. For example, the following are some of the variables that are automatically customized (see Appendix B, “Appendix B. Sample kscfg.tmpl File” on page 27 for all of the variables):

#MGMTSVR_HOSTNAME#	Replaced with the host name of the management server.
#NFS_HOSTNAME#	Replaced with the host name of the management server.
#NFS_DIR#	Replaced with the directory on the management server that contains the Red Hat installation images; for example, **/tftpboot/rh71**.

The following information is in the **kscfg.tmpl** file and in the generated Kickstart configuration file for each node, and can be modified if needed:

- Initial root password (defaults to ″cluster″)
- Number of disks installed
  - Type of each disk
  - Size of each disk

  - Whether to clear partitions
- Disk partition table

  A sample disk partition table is provided, which can be modified.
- RPMs to install

  A standard list is provided; you can modify the list, or you can use your own list.
- Language and time zone
- Post-install script

  This script does standard set-up and CSM-specific set-up. The script may be modified to suit your installation provided the marked sections are not altered.

Note: See the annotated **kscfg.tmpl** sample file in the Appendix, and the **setupks** and **kscfg.tmpl** man pages or the *IBM Cluster Systems Management for Linux Technical Reference* for further details.

Data is collected for each node for the **/etc/dhcp.conf** file and Kickstart configuration file.

The following information is collected for the **/etc/dhcp.conf** file. Any existing **/etc/dhcp.conf** file is saved to **/etc/dhcp.conf.orig** and a new file is created. After running **setupks**, any entries that were previously in **/etc/dhcp.conf** should be replaced manually.

- MAC address

  This information is collected by the **getmacs** command and goes into the *Macaddr* attribute of the PreManaged node object and in the **/etc/dhcp.conf** file, provided a value does not already exist. An existing value will not be replaced. See *IBM Cluster Systems Management for Linux Technical Reference* for details.
- IP address of the node
- Net mask (defaults to the net mask of the management server or can be specified on the **setupks** command line)
- Gateway (defaults to the gateway of the management server or can be specified on the **setupks** command line)
- Nameservers (defaults to the nameservers of the management server or can be specified on the **setupks** command line; is also used in the **kscfg.tmpl** file)

A **firstboot** file is also generated for each node. The template file is in **/opt/csm/install/firstboot.tmpl** and the generated files for each node are in **/tftpboot/bin/<node-hostname>.firstboot**. The **firstboot** script is run the first time a node boots from its hard drive after it is installed. The purpose of the **firstboot** script is to run the **makenode** command. Code may be added to the template or script to run additional functions during **firstboot**.

Running installnode

This command is used to install the nodes in the cluster by running **makenode** remotely on each node for a CSM-only installation or by using Kickstart to run **makenode** for a full installation. 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 **monitorinstall** command displays the installation status for each node. In addition, the **makenode.log** file is maintained on each node in **/var/log/csm/makenode.log** that contains information on what happened during installation on each node and a similar **installnode.log** file is maintained on the management server with information on the whole installation process. For a full installation, a **/var/log/csm/kickstart.log** file and **/var/log/csm/firstboot.log** file on each node contain information about Kickstart and firstboot post-installation.

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. For a CSM-only installation, 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 **/tftpboot** 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 **definenode** commands must have been run successfully for a CSM-only installation, For a full installation, the **setupks** command also must have been run successfully.
4. Verify that the node definitions are correct, or change the node definitions. See “Verifying node definitions” on page 9. When you are satisfied with the node definitions, run **installnode**.

Run installnode

1. Run **installnode**. Note that the **installnode** command runs synchronously for a CSM-only installation; that is, when **installnode** completes, the installation of the cluster is complete. But for a full installation, **installnode** runs asynchronously; that is, immediately after the installation process is initiated (the node is rebooted), **installnode** exits even though the installation may not be complete.
2. After you have a working cluster, see “Getting started with the newly installed cluster” on page 13 to make sure that the cluster is functioning successfully. (If you are performing a CSM-only installation, you should have a working cluster after running **installnode**.) Run the **monitorinstall** command to see how the installation is progressing.
3. See **/var/log/csm/installnode.log** on the management server or **/var/log/csm/makenode.log** on the managed node for information on what happened on each node during installation if there is a problem. To determine how to handle a problem, see “Chapter 4. FAQs, hints, tips, and troubleshooting” on page 21.
4. For a full installation, see **/var/log/csm/kickstart.log** and **firstboot.log** for information on the processing of the Kickstart post-installation script and the firstboot script.

After a full install, the boot order in the BIOS of the node can be left as is: floppy, CD-ROM, network, hard disk. In this case, every time the node boots it uses **dhcp** to contact the management server, which uses **pxelinux** to boot the node from its hard drive. Or after the full installation is complete, you can change the boot order in the BIOS back to: floppy, DC-ROM, hard disk, network.

Check installation status

The installation monitor tool is started by running the **monitorinstall** command. The tool displays the status of the installation on each of the nodes. It returns the following types of status: installed, installing, not installed, failed install.

You can also run the **rconsole** command to view installation progress on each node for a full installation. Type the following:

```
rconsole -Pa
```

Adding a node to an existing cluster

You can add another node to the cluster either by running **definenode**, **setupks** if this is a full installation, 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 -Al
```

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

```
lsnode -AlP
```

2. To define the node, type:

```
definenode
```

You will be prompted for the required information.

3. After **definenode** is completed successfully, run **setupks** if this is a full installation; then run **installnode**.
4. Run **monitorinstall** to check on the status of the installation process.
5. To verify that the node is installed, type:

```
lsnode -Al nodename
```

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 **definenode**, **setupks** if this is a full installation, and **installnode** 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 begin working with the Configuration File Manager, use the following example. The example sets up the **cforce** command to run whenever the **/cfmroot/etc/passwd** file has changed, distributing the changed file across all nodes in the cluster.

  ```
  mkdir /cfmroot/etc
  cp /etc/passwd /cfmroot/etc/passwd
  startcondresp "CFMRootModTimeChanged" "CForce"
  ```
- 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 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
  ```

- To set up node groups called **servers** and **admin** and then run the **vmstat** command on them, do the following:

  1. To form a node group called **servers** that consists of nodes c5bn07, c5bn08, c5bn09, c5bn10, and c5bn11, type:

     ```
     nodegrp -a c5bn07,c5bn08,c5bn09,c5bn10,c5bn11 servers
     ```

  2. To form a node group called **admin** that consists of nodes c5bn12 and, c5bn13, type:

     ```
     nodegrp -a c5bn12,c5bn13 admin
     ```

  3. To display the existing node groups, type:

     ```
     nodegrp -l
     ```

     The output is similar to:

     ```
     admin
     servers
     ```

  4. To display the nodes in the node group **servers**, type:

     ```
     nodegrp servers
     ```

     The output is:

     ```
     c5bn07.ppd.pok.ibm.com
     c5bn08.ppd.pok.ibm.com
     c5bn09.ppd.pok.ibm.com
     c5bn10.ppd.pok.ibm.com
     c5bn11.ppd.pok.ibm.com
     ```

  5. To run the **vmstat** command on the node group **servers** and have formatted output, type:

     ```
     dsh -N servers vmstat | dshbak
     ```

     The output is similar to:

     ```
     HOST: c5bn08.ppd.pok.ibm.com
     ----------------------------
        procs                      memory    swap          io     system         cpu
      r  b  w   swpd   free   buff  cache  si  so    bi    bo   in    cs  us  sy  id
      0  4  1 442440 192576  56292 635808   0   0     0     0    1     1   0   0   0

     HOST: c5bn09.ppd.pok.ibm.com
     ----------------------------
        procs                      memory    swap          io     system         cpu
      r  b  w   swpd   free   buff  cache  si  so    bi    bo   in    cs  us  sy  id
      0  4  1 423692 214232  56240 615396   0   0     0     0    1     1   0   0   0

     HOST: c5bn10.ppd.pok.ibm.com
     ----------------------------
        procs                      memory    swap          io     system         cpu
      r  b  w   swpd   free   buff  cache  si  so    bi    bo   in    cs  us  sy  id
      0  4  1 405904 162404  56248 604424   0   0     0     0    4     1   0   0   1
     ```

```
HOST: c5bn11.ppd.pok.ibm.com
----------------------------
   procs                      memory    swap          io     system         cpu
 r  b  w   swpd   free   buff  cache  si  so    bi    bo   in    cs  us  sy  id
 0  4  1 443564 135240  56212 636256   0   0     0     0    4     1   0   0   1
```

Chapter 2. Backing up and restoring CSM data

The chapter discusses backing up and restoring CSM data.

Backing up CSM data

The following directories and all their subdirectories should be backed up periodically. These directories are used by Cluster Systems Management to store persistent data and other configuration information. In addition, if major configuration changes are made by means of CSM commands, a back up should be performed.

These directories only need to be backed up on the management server.

- CSM Configuration Data: **/etc/opt/csm**
- CSM Runtime Data: **/var/opt/csm**
- RMC Persistent Data: **/var/ct** (See “Backing up and restoring /var/ct”.)

Restoring CSM data

Do the following procedure to restore CSM data:

1. Reinstall CSM if necessary.
2. Copy **/etc/opt/csm** and **/var/opt/csm** from the backup media.
3. Refresh RMC. (See “Backing up and restoring /var/ct”.):

Backing up and restoring /var/ct

To perform a backup, do the following steps:

1. Run the following command to stop the RMC subsystems.

   ```
   /usr/sbin/rsct/bin/rmcctrl -z
   ```

   The command does not complete until all RMC subsystems are stopped.
2. Perform a backup of the **/var/ct** directory
3. Run the following command to restart the RMC subsystems.

   ```
   /usr/sbin/rsct/bin/rmcctrl -s
   ```

To restore the **/var/ct** directory, do the following steps:

1. Stop the RMC subsystems by running the following command:

   ```
   /usr/sbin/rsct/bin/rmcctrl -z
   ```

   The command does not complete until all RMC subsystems are stopped.
2. If in HA cluster mode, run the following command to stop the cluster subsystems:

   ```
   /usr/sbin/rsct/bin/cthactrl -k
   ```
3. Restore the **/var/ct** directory.
4. If in HA cluster mode, run the following command to start the cluster subsystems:

   ```
   /usr/sbin/rsct/bin/cthactrl -s
   ```
5. To restart the RMC subsystems, run the following command

   ```
   /usr/sbin/rsct/bin/rmcctrl -s
   ```

Chapter 3. Uninstalling CSM

At the time of this writing, the **rpm** command does not reorder packages based on dependencies when doing an uninstall. Instead, the packages are uninstalled in the order they are specified on the command line. To properly uninstall the csm and rsct packages and the packages on which they depend, the packages should be specified as follows:

On the management server:

```
rpm -e csm.server csm.dsh csm.core rsct.core rsct.core.utils src
```

On a node:

```
rpm -e csm.client csm.dsh csm.core rsct.core rsct.core.utils src
```

If packages that depend on the csm and rsct packages are uninstalled with the same invocation of the **rpm** command, those packages should be specified before the csm and rsct packages.

Fixing the ACL file after uninstalling CSM

When csm.server or csm.client is uninstalled, they do not clean up their stanzas in the RMC ACL file (**/var/ct/cfg/ctrmc.acls**). If rsct.core is also still installed, this causes RMC to ignore the ACL file and restrict access to local root only. The following stanzas (if present) should be manually deleted (along with all entries they contain) from **/var/ct/cfg/ctrmc.acls** if either csm.server or csm.client is uninstalled and rsct.core is still installed:

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

After the ACL file has been modified, run the following command:

```
/usr/bin/refresh -s ctrmc
```

Chapter 4. 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 Administration Guide*.

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 How do I set up **dsh**?

Answer If you are using **rsh**, make sure each node has a **/root/.rhosts** directory that contains the management server host name. Adding ″rsh″ to the end of **/etc/securettty** on each node allows root to **rsh** into the node. Also make sure the **rsh** daemon is running on each node: use `chkconfig --list rsh` to check it, or `chkconfig rsh on` to turn on **rsh**.

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

Answer Consult the **/var/log/csm/installnode.log** file on the management server and **/var/log/csm/makenode.log** on the nodes. For a full installation, run `monitorinstall -l` to see detailed status of the install, or see the **kickstart.log** or **firstboot.log** file on the nodes.

Question What do I do if I have one or more nodes that are displayed as ManagedNodes by the **lsnode** command and as PreManagedNodes by the **lsnode -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 skip the MAC address collection during **setupks**?

Answer Assign a MAC address to each node by running `chnode <node> -a "Macaddr=00:01:02:03:04:05"` for each node.

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 /opt/bin/mgmtsvr <newhostname>
```

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 Can I install more than one management server?

Answer Currently, only one management server is allowed per CSM cluster. In the future, more than one management server is planned for high availability.

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.

Question Does deleting a node from the cluster also uninstall the CSM RPMs from the node?

Answer No, the CSM software is left on the node. To uninstall the CSM RPMs, type:

```
rpm -e csm.client csm.dsh csm.core rsct.core rsct.core.utils src
```

Question If I want to pass in a null value for a hardware control point or a console server when I am running the **addnode** command, how do I do it correctly so that I am not prompted for the value later?

Answer To pass in a null value to the **definenode** command and avoid being prompted later, use quotation marks as in the following example:

```
definenode -n node1 -c 1 -C" " -H" " -t netfinity
```

To do the same thing with the **addnode** command is a bit more complex. The syntax that must be used in the addnode command is as follows:

```
addnode -n node1 -c 1 -C\"\" -H\"\" -t netfinity
```

See the man pages for more details.

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 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. 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 A. Node Attributes Template

See *IBM Cluster Systems Management for Linux Remote Control Guide and Reference* for an example of a filled-out node-attributes template. Note that the console port number is the physical port to which the serial port of the node is connected in the console server hardware. Duplicate the template and complete it before you install CSM.

Table 1. Node Attributes Table: Template

Hostname	*HWControlPoint*	*Power Method*	*Svc ProcName*	*ConsoleServerName*	*Console Server Number*	*Console Method*	*Console PortNum*	*HWType*	*InstallMethod*

Appendix B. Sample kscfg.tmpl File

The following is a sample **kscfg.tmpl File**. See the **kscfg.tmpl File** man page or the *IBM Cluster Systems Management for Linux Technical Reference* for more information.

```
#
#Kickstart Configuration File for RedHat 7.1 Only
#

lang en_US
keyboard "us"

network --bootproto dhcp

#
# Where's the source?
# nfs --server hostname.of.server or IP --dir /path/to/RH/CD/image
#
nfs --server #NFS_HOSTNAME# --dir #NFS_DIR#


#
# Clear the MBR
#
zerombr yes

#
# Wipe out the disk
#
clearpart --all

#
# Disk partition table. Customize it to fit your needs.
# /boot is strongly recommended
#
part /boot --size 50   --ondisk sda
part /     --size 1024 --ondisk sda
part swap  --size 1024 --ondisk sda
part /var  --size 1024 --ondisk sda
part /opt  --size 512  --ondisk sda
part /usr  --size 2048 --ondisk sda

#
# install or upgrade
#
install

#
# mouse selection
#
#mouse genericps/2 --emulthree
mouse none

#
# Select a time zone
# Add the --utc switch if your hardware clock is set to GMT
#
#timezone US/Hawaii
#timezone US/Pacific
#timezone US/Mountain
#timezone US/Central
#timezone US/Eastern
timezone --utc US/Eastern

#
```

```
# Don't do X
#
skipx

#
# To generate an encrypted root password use:
#
# perl -e 'print crypt("blah","Xa") . "\n";'p
# openssl passwd -apr1 -salt xxxxxxxx password
#
# where "blah" is your root password.
#
# --iscrypted is broken
#
#rootpw --iscrypted XaLGAVe1C41x2
#rootpw XaLGAVe1C41x2 --iscrypted
rootpw cluster

auth --useshadow --enablemd5

#
# Lilo
#
lilo --location mbr

#
# Reboot after installation
#
reboot

#
#end of section
#

#
# Choose the packages and package groups to install
#
%packages
@ Networked Workstation
compat-libstdc++
kernel-smp
kernel-headers
glibc-devel
binutils
cpp
gcc
flex
autofs
pdksh
tcsh
ntp
xosview
XFree86
tftp
xinetd
rsh
rsh-server
telnet-server
pvm
pvm-gui
psacct
nfs-utils
raidtools
ucd-snmp
rsync
yp-tools
ypserv
```

```
ypbind
m4
sendmail-cf
%post
#
# Post Install script
#
# This script may be modified as needed.  However, certain sections should
# not be altered or else the CSM installation will not work correctly. These
# sections are marked with the text "DO NOT ERASE THIS SECTION".
#

#
# The LOGGER variable is used to send items to some kind of log.  It can be
# set to 'write_log' in order to call the 'write_log' function (defined below)
# which sends the log output to /var/log/csm/kickstart.log on the node.
# If LOGGER is set to /usr/bin/logger, the Linux syslog facility is used
# instead.
#
#export LOGGER=/usr/bin/logger
export LOGGER=write_log

#
# Set variables
#
# The following variables are substituted into the kickstart template file
# from the setupks tool
#
# MGMTSVR_HOSTNAME = management server (Long Hostname)
# MGMTSVR_SHORTHOST = management server (Short Hostname)
# MGMTSVR_IP = management server (IP)
# NODE_HOSTNAME = node hostname (Long Hostname)
# NFS_HOSTNAME = NFS Server (Long Hostname) - same as MGMTSVR_HOSTNAME
# NFS_IP = NFS Server (IP) - same as MGMTSVR_IP
# NFS_DIR = Path to RedHat CD image on NFS Server
# TFTP_IP = TFTP Server (IP) - same as MGMTSVR_IP
# NAMESERVERS = list of comma delimited name servers(IPs or Hostnames)
# DNSDOMAIN = DNS domain name of the node(NAME)
# SETUP_RSH = Set to "Y" if $DSH_REMOTE_CMD variable on the
#   management server is set to "rsh" or is blank.
#
export KICKSTART_LOGFILE=/var/log/csm/kickstart.log
export MGMTSVR_HOSTNAME=#MGMTSVR_HOSTNAME#
export MGMTSVR_SHORTHOST=#MGMTSVR_SHORTHOST#
export MGMTSVR_IP=#MGMTSVR_IP#
export NODE_HOSTNAME=#NODE_HOSTNAME#
export NFS_HOSTNAME=#NFS_HOSTNAME#
export NFS_IP=#NFS_IP#
export TFTP_IP=#TFTP_IP#
export NAMESERVERS=#NAMESERVERS#
export DNSDOMAIN=#DNSDOMAIN#
export SETUP_RSH=#SETUP_RSH#

#
# Hardcode some other variables
#
export GATEWAY=$(netstat -rn | grep 0.0.0.0 | tail -1 | awk '{print $2}')
export CSM_MANPATH=/opt/man
export CSM_PATH=/opt/bin
export TFTPBOOT=/tftpboot
export ATFTP=$TFTPBOOT/bin/atftp

#
# Function to add some text to the kickstart.log file on the node.
#
write_log()
{
```

```
    mkdir -p /var/log/csm
    echo "'date': $*" >> $KICKSTART_LOGFILE
} # END write_log

#
# Use atftp to get a file from the management server.  The remote file is
# relative to /tftpboot
#
atftp_get()
{
    _remote_file=$1
    _local_file=$2

    # Use this format if LOGGER=/usr/bin/logger
    #echo "get $_local_file $remote_file" | $ATFTP $TFTP_IP 2>&1 | $LOGGER

    # Use this format if LOGGER=write_log
    $LOGGER "'echo \"get $_remote_file $_local_file\" | $ATFTP $TFTP_IP 2>&1'"
} # END atftp_get

#
# Use atftp to put a file to the management server.  The remote file is
# relative to /tftpboot
#
atftp_put()
{
    _local_file=$1
    _remote_file=$2

    # Use this format if LOGGER=/usr/bin/logger
    #echo "put $_remote_file $local_file" | $ATFTP $TFTP_IP 2>&1 | $LOGGER

    # Use this format if LOGGER=write_log
    $LOGGER "'echo \"put $_local_file $_remote_file\" | $ATFTP $TFTP_IP 2>&1'"
} # END atftp_put

#
# Function to add a line to the status file on the management server.
# The status file is in the /tftpboot/status directory on the management server.
#
# The first argument to this function is the message.  It should be quoted
# if it contains spaces.  The second optional argument is a status.  This is
# typically used to send the exit code of a command to the status log.
#
# The format of the output message in the status file is:
#   <date>: <message>: <status>
# If <status> is not provided, the format of the message in the status file is:
#    <date>: <message>
#
# This function uses atftp to transfer the status file back and forth.  The
# atftp server should have been setup on the management server automatically
# when csm.server was installed.  The atftp client resides in the mounted
# /tftpboot filesystem.  So this function cannot be used until after the
# /tftpboot filesystem is mounted.
#
write_status()
{
    msg=$1;
    status=${2:-''}

    remote_file="status/$NODE_HOSTNAME"
    local_file="/tmp/${NODE_HOSTNAME}.status"

    $LOGGER "Writing status to $remote_file.  MSG=$msg"

    # Get the previous version of the status file
    /bin/rm -f $local_file
```

```
    #$LOGGER "'echo \"get $remote_file $local_file\" | $ATFTP $TFTP_IP 2>&1'"
    atftp_get $remote_file $local_file

    # Add a line to the status file
    if [ "$status" ]
    then
echo "'date': $msg: status=$rc" >> $local_file
    else
echo "'date': $msg" >> $local_file
    fi

    # Write the new version of the status file
    #$LOGGER "'echo \"put $local_file $remote_file\" | $ATFTP $TFTP_IP 2>&1'"
    atftp_put $local_file $remote_file

    /bin/rm -f $local_file
}  # END write_status


$LOGGER "Starting Install"

#netstat -rn | $LOGGER # Use this format if LOGGER=/usr/bin/logger
$LOGGER "'netstat -rn'" # Use this format if LOGGER=write_log

#
# Setup Syslog
#
# Turn syslog so we can use $LOGGER to send messages to syslogd running on a
# management node. Use $LOGGER.
#
# Uncomment this if you want all syslog messages from the node to be sent to
# the management server.  Be sure to also uncomment the section that says
# "Setup Syslog again" (below)
#
#mv /etc/syslog.conf /etc/syslog.conf.ORIG
#echo "*.* @$MGMTSVR_IP" >/etc/syslog.conf
#/etc/rc.d/init.d/syslog start
#$LOGGER "Install: syslog setup"

###################### DO NOT ERASE THIS SECTION (begin) ###################
#
# Mount /tftpboot
#
$LOGGER "Mounting $TFTPBOOT"
# Use this format if LOGGER=/usr/bin/logger
#mkdir $TFTPBOOT 2>&1 | $LOGGER
#mount -o nolock $NFS_IP:$TFTPBOOT $TFTPBOOT 2>&1 | $LOGGER
# Use this format if LOGGER=write_log
$LOGGER "'mkdir $TFTPBOOT 2>&1'"
$LOGGER "'mount -o nolock $NFS_IP:$TFTPBOOT $TFTPBOOT 2>&1'"

write_status "Kickstart RPM Installation Complete."
write_status "Starting Kickstart Post-Install."

###################### DO NOT ERASE THIS SECTION (end) ######################

#
# Hardcode eth0 (optional, recommended)
#
$LOGGER "Setting up eth0"
IP0=$(ifconfig eth0 | grep inet | awk '{print $2}' | awk -F: '{print $2}')
BC0=$(ifconfig eth0 | grep inet | awk '{print $3}' | awk -F: '{print $2}')
SM0=$(ifconfig eth0 | grep inet | awk '{print $4}' | awk -F: '{print $2}')
cd /etc/sysconfig/network-scripts
cp ifcfg-eth0 ifcfg-eth0.ORIG
echo "DEVICE=eth0
BOOTPROTO=none
```

```
ONBOOT=yes
USERCTL=no
IPADDR=$IP0
BROADCAST=$BC0
NETMASK=$SM0" >ifcfg-eth0

#
# Setup /etc/resolv.conf
#
$LOGGER "Setting up /etc/resolv.conf"
mv /etc/resolv.conf /etc/resolv.conf.ORIG
if [ "$DNSDOMAIN" ]
then
    echo "search $DNSDOMAIN" >/etc/resolv.conf
fi
for i in $(echo $NAMESERVERS | tr ',' ' ')
do
echo "nameserver $i"
done >>/etc/resolv.conf

#
# Setup /etc/sysconfig/network
#
$LOGGER "Setting up /etc/sysconfig/network"
HOSTNAME=$(host $IP0 2>/dev/null | awk '{print $5}' | sed 's/\.$//')
SHORT_HOSTNAME=$(echo $HOSTNAME | awk -F. '{print $1}')
cp /etc/sysconfig/network /etc/sysconfig/network.ORIG
echo "NETWORKING=yes
HOSTNAME=$HOSTNAME
FORWARD_IPV4=yes
GATEWAYDEV=eth0
GATEWAY=$GATEWAY" >/etc/sysconfig/network

#
# Setup /etc/hosts
# Add myself and the management server
#
$LOGGER "Setting up /etc/hosts"
echo "127.0.0.1 localhost

$IP0 $SHORT_HOSTNAME $HOSTNAME
$MGMTSVR_IP $MGMTSVR_SHORTHOST $MGMTSVR_HOSTNAME" > /etc/hosts

#
# Add users (optional)
#
# The following line produces an encrypted password of "cluster"
# perl -e 'print crypt("cluster","Xa") . "\n";'p
#
$LOGGER "Adding users"
adduser admin -u 1000 -p 'perl -e 'print crypt("cluster","Xa") . "\n";''

#
# Enable rsh, add management server to /root/.rhosts and add 'rsh' to
# /etc/securetty to allow root to rsh to the node.
#
if [ "${SETUP_RSH}" = "Y" ]
then
if [ -r /etc/xinetd.d/rsh ]
then
$LOGGER "Enable rsh"
# Could use chkconfig to do this
#perl -pi -e 's/^(\tdisable[^=]*=) yes/$1 no/' /etc/xinetd.d/rsh
chkconfig rsh on

$LOGGER "Allow root to rsh"
echo "rsh" >> /etc/securetty
```

```
$LOGGER "Setting up /root/.rhosts"
cd /root
echo "$MGMTSVR_SHORTHOST" >.rhosts
echo "$MGMTSVR_HOSTNAME" >>.rhosts
chmod 600 .rhosts
fi
fi

#
# Setup PAM (optional, recommended)
#
# Uncomment this if you want to enable ftp, login, rlogin, rsh, sshd
#$LOGGER "Setup PAM"
#cd /etc/pam.d
#for i in ftp login rlogin rsh sshd
#do
# echo "account    required     /lib/security/pam_access.so" >>$i
#done
#echo "-:ALL EXCEPT root:ALL" >>/etc/security/access.conf
#cp /etc/security/access.conf /etc/security/access.conf.BOOT
#echo "cp /etc/security/access.conf.BOOT /etc/security/access.conf >/dev/null 2>&1" >>/etc/rc.d/rc.local


#
# Setup man paths in /etc/man.config.  Add CSM_MANPATH after /usr/local/man.
#
$LOGGER "Setup man paths"
cp -f /etc/man.config /etc/man.config.ORIG
FIXEDMANPATH=$(echo $CSM_MANPATH | sed 's/\//\\\//g')
perl -pi -e "s/^MANPATH[\t ]+\/usr\/local\/man$/MANPATH\t\/usr\/local\/man\nMANPATH $FIXEDMANPATH/" /etc/man.config

#
# Setup paths.
#  Add PATH definitions to /etc/profile.d/csm.sh and /etc/profile.d/csm.csh
#  Add section to /etc/csh.login that source's all the *.csh files
#     in /etc/profile.d/  (just like /etc/profile does for *.sh files).
#
$LOGGER "Setup PATHS"
echo "export PATH=\$PATH:$CSM_PATH" >/etc/profile.d/csm.sh
echo "setenv PATH \"\${PATH}:$CSM_PATH\"" >/etc/profile.d/csm.csh
chmod 755 /etc/profile.d/csm.*
echo "
test -d /etc/profile.d
if (\$status == 0) then
set nonomatch
foreach i ( /etc/profile.d/*.csh )
test -f \$i
if (\$status == 0) then
source \$i
endif
end
unset i nonomatch
endif" >>/etc/csh.login

#
# Setup services
#
$LOGGER "Setup services"
#chkconfig --del apmd
#chkconfig --del gpm
#chkconfig --del kudzu
#chkconfig --del lpd
#chkconfig --del pcmcia
#chkconfig --del linuxconf
#chkconfig --del sendmail
#chkconfig --del xfs
```

```
#chkconfig --del httpd
#chkconfig --del autofs
#chkconfig --del identd
#chkconfig --del isdn
#chkconfig --del pppoe
#chkconfig --level 345 ntpd on
#chkconfig --level 345 sshd on
#chkconfig --level 345 snmpd on
chkconfig telnet on

#
# Setup Syslog Again
#
# Uncomment this if you want all syslog messages from the node to be sent to
# the management server.  Be sure to also uncomment the section that says
# "Setup Syslog" (above)
#
#$LOGGER "Syslog setup again"
#echo "*.* @$MGMTSVR_IP" >/etc/syslog.conf

#
# Set local and remote complete install flag
# Put file in /dev/sda1
date >/boot/install_complete

###################### DO NOT ERASE THIS SECTION (begin) ###################
#
# Create a pxelinux configuration file for this node and tftp it to the
# management server.  This pxelinux file will cause the node to boot to
# the hard disk on the next reboot after the install is complete.  The
# CSM code will be installed during the first reboot of the hard disk.
#
$LOGGER "Update pxelinux configuration on management server"
IP=$(ifconfig eth0 | grep inet | awk '{print $2}' | awk -F: '{print $2}')
HEX=$(
for i in $(echo $IP | tr '.' ' ')
do
printf "%02x" $i
done | tr '[a-z]' '[A-Z]'
)
cd /tmp

#
# Use atftp to copy the HEX file over to the management server.
# The HEX file contains a new pxelinux configuration that causes the node
# to boot from hard disk on the next reboot.
#
echo "#boot" >$HEX
#$LOGGER "'echo \"get pxelinux.cfg/$HEX blah\" | $ATFTP $TFTP_IP 2>&1'"
atftp_get pxelinux.cfg/$HEX blah

grep SERIAL blah >>$HEX
echo "DEFAULT xCAT
LABEL xCAT
LOCALBOOT 0" >>$HEX

#$LOGGER "'echo \"put $HEX pxelinux.cfg/$HEX\" | $ATFTP $TFTP_IP 2>&1'"
atftp_put $HEX pxelinux.cfg/$HEX

#
# Copy the node's firstboot script from /tftpboot to the node
#
mkdir -p /opt/csm/bin/
/bin/cp /tftpboot/bin/$NODE_HOSTNAME.firstboot /opt/csm/bin/firstboot
chmod 755 /opt/csm/bin/firstboot

#
```

```
# Add the firstboot script to /etc/inittab.  The firstboot script runs only during
# the first reboot of the hard disk.  It runs 'makenode' to install CSM on the node.
#
if grep "^csm:" /etc/inittab > /dev/null
then
# Line already exists so leave it
true
else
$LOGGER "Adding firstboot entry to /etc/inittab"
# Add comment and firstboot entry to /etc/inittab
echo "
#CSMFIRSTBOOT: The firstboot script runs only after the node is booted the
#CSMFIRSTBOOT: first time after full install.  Firstboot then removes this
#CSMFIRSTBOOT: entry from /etc/inittab.
csm:345:wait:/opt/csm/bin/firstboot" >> /etc/inittab
fi

#
# Copy the kickstart log to the management server in /tftpboot/status
#
local_file=$KICKSTART_LOGFILE
remote_file="status/$NODE_HOSTNAME.kickstart.log"
atftp_put $local_file $remote_file

#
# Tell the management server that the install is complete.
#
write_status "Kickstart Post-Install Complete."
write_status "Rebooting to hard disk."

#
# Unmount $TFTPBOOT
#
$LOGGER "Unmounting $TFTPBOOT"
cd /
umount $TFTPBOOT
rmdir $TFTPBOOT
###################### DO NOT ERASE THIS SECTION (end) ######################

cd /
$LOGGER "Install complete"
exit 0
```

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 Road
Poughkeepsie, New York 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 terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

- IBM, AIX, AIX 5L, Netfinity, RS/6000, SP, and the e (logo) are trademarks or registered trademarks of International Business Machines Corporation.
- Equinox is a trademark of Equinox Systems, Inc.
- Linux is a registered trademark of Linus Torvalds.
- Myrinet is a trademark of Myricom, Inc.
- Red Hat and RPM are trademarks of Red Hat, Incorporated.
- Java and all Java-based trademarks and logos are registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
- UNIX is a registered trademark of The Open Group in the United States and other countries.
- Other company, product, or 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:

Advanced TFTP	Advanced TFTP (atftp) is a client/server implementation of the TFTP protocol. It is licensed under the GNU General Public License (GNU GPL), which can be found at **http://www.gnu.org/copyleft/gpl.html**.
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 or ″reverse-telnet″ to console servers or both.
DBD-CSV, Text-CSV_XS	Licensed by GPL or Artistic, these are dynamically loaded Perl modules.
DBI	Licensed by GPL or Artistic, this is a dynamically loaded Perl module.
fping	Copyrighted by Stanford University, this is executed as a separate binary.
Kerberos	Provides authentication of the execution of remote commands.
Perl	Practical Extraction and Report Language is licensed under the Artistic license.
Perl-to-C extensions	Practical Extraction and Report Language-to-C extensions is distributed 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.
SYSLinux	SYSLinux includes PXELINUX, which CSM uses to control the behavior of network boots. SYSLinux is licensed under the GNU GPL.

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 and license conditions in the documentation must be respected. You can find version and distribution information for each of these products in the *Specified Operating Environment* section of this book. For these non-IBM products, the following license terms apply in lieu of the International Program License Agreement.

The inclusion herein of copies of various licenses is not meant to imply endorsement of the principles, methodologies, or views that are contained therein, either express or implied.

For the fping open source code, see the following URL: **http://www.fping.com** (The license is imbedded in the source code.)

For the fping open source code, the following terms apply:

```
Copyright (c) 1991, 1994, 1997 Board of Trustees
              Leland Stanford Jr. University
```

Redistribution and use in source and binary forms are permitted provided that the above copyright notice and this paragraph are duplicated in all such forms and that any documentation, advertising materials, and other materials related to such distribution and use acknowledge that the software was developed by Stanford University. The name of the University may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED ″AS IS″ AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTIBILITY AND FITNESS FOR A PARTICULAR PURPOSE.

For the DBD-CSF, DBI, Text-CSV_XS and SQL-Statement open source code licensed under the ″Artistic License,″ the following terms apply:

```
Definitions

"Package" refers to the collection of files distributed by the
Copyright Holder, and derivatives of that collection of files created
```

```
through textual modification.

"Standard Version" refers to such a Package if it has not been
modified, or has been modified in accordance with the wishes of the
Copyright Holder as specified below.

"Copyright Holder" is whoever is named in the copyright or copyrights
for the package.

"You" is you, if you're thinking about copying or distributing this
Package.

"Reasonable copying fee" is whatever you can justify on the basis of
media cost, duplication charges, time of people involved, and so on.
(You will not be required to justify it to the Copyright Holder, but
only to the computing community at large as a market that must bear the
fee.)

"Freely Available" means that no fee is charged for the item itself,
though there may be fees involved in handling the
item. It also means that recipients of the item may redistribute
it under the same conditions they received it.
```

1. You may make and give away verbatim copies of the source form of the Standard Version of this Package without restriction, provided that all copies contain the original copyright notices and associated disclaimers.
2. You may apply bug fixes, portability fixes and other modifications derived from the Public Domain or from the Copyright Holder. A Package modified in such a way shall still be considered the Standard Version.
3. You may otherwise modify your copy of this Package in any way, provided that you insert a prominent notice in each changed file stating how and when you changed that file, and provided that you do at least ONE of the following:
   a. place your modifications in the Public Domain or otherwise make them Freely Available, such as by posting said modifications to Usenet or an equivalent medium, or placing the modifications on a major archive site such as uunet.uu.net, or by allowing the Copyright Holder to include your modifications in the Standard Version of the Package.
   b. use the modified Package only within your corporation or organization.
   c. rename any non-standard executables so the names do not conflict with standard executables, which must also be provided, and provide a separate manual page for each non-standard executable that clearly documents how it differs from the Standard Version.
   d. make other distribution arrangements with the Copyright Holder.
4. You may distribute the programs of this Package in object code or executable form, provided that you do at least ONE of the following:
   a. distribute a Standard Version of the executables and library files, together with instructions (in the manual page or equivalent) on where to get the Standard Version.
   b. accompany the distribution with the machine-readable source of the Package with your modifications.
   c. give non-standard executables non-standard names, and clearly document the differences in manual pages (or equivalent), together with instructions on where to get the Standard Version.
   d. make other distribution arrangements with the Copyright Holder.
5. You may charge a reasonable copying fee for any distribution of this Package. You may charge any fee you choose for support of this Package. You may not charge a fee for this Package itself. However, you may distribute this Package in aggregate with other (possibly commercial) programs as part of a larger (possibly commercial) software distribution provided that you do not advertise this Package as a product of your own. You may embed this Package's interpreter within an executable of yours (by linking); this shall be construed as a mere form of aggregation, provided that the complete Standard Version of the interpreter is so embedded.

6. The scripts and library files supplied as input to or produced as output from the programs of this Package do not automatically fall under the copyright of this Package, but belong to whomever generated them, and may be sold commercially, and may be aggregated with this Package. If such scripts or library files are aggregated with this Package via the so-called ″undump″ or ″unexec″ methods of producing a binary executable image, then distribution of such an image shall neither be construed as a distribution of this Package nor shall it fall under the restrictions of Paragraphs 3 and 4, provided that you do not represent such an executable image as a Standard Version of this Package.
7. C subroutines (or comparably compiled subroutines in other languages supplied by you and linked into this Package in order to emulate subroutines and variables of the language defined by this Package shall not be considered part of this Package, but are the equivalent of input as in Paragraph 6, provided these subroutines do not change the language in any way that would cause it to fail the regression tests for the language.
8. Aggregation of this Package with a commercial distribution is always permitted provided that the use of this Package is embedded; that is, when no overt attempt is made to make this Package's interfaces visible to the end user of the commercial distribution. Such use shall not be construed as a distribution of this Package.
9. The name of the Copyright Holder may not be used to endorse or promote products derived from this software without specific prior written permission.
10. THIS PACKAGE IS PROVIDED ″AS IS″ AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTIBILITY AND FITNESS FOR A PARTICULAR PURPOSE.

For the Advanced TFTP (atftp), cfengine, and SYSLinux code, licensed under the GNU GENERAL PUBLIC LICENSE Version 2, June 1991, see the following.

```
Copyright (c) 1989, 1991 Free Software Foundation, Inc. 59 Temple
Place, Suite 330, Boston, MA 02111-1307 USA
```

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The ″Program″, below, refers to any such program or work, and a ″work based on the Program″ means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term ″modification″.) Each licensee is addressed as ″you″.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program. You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.
2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:
   a. You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
   b. You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
   c. If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

   These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:
   a. Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
   b. Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
   c. Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

   The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

   If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.
4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.
5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.
6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.
7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the

Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances. It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.
9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

   Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and ″any later version″, you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.
10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM ″AS IS″ WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.
12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

Index

Special Characters

A

C

D

E

F

G

H

I

K

L

M

N

O

P

R

S

T

V

W

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

IBM Cluster Systems Management for Linux®
Planning and Installation Guide
Version 1.1

Publication No. SA22-7853-02

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

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	☐	☐	☐	☐	☐

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

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	☐	☐	☐	☐	☐
Complete	☐	☐	☐	☐	☐
Easy to find	☐	☐	☐	☐	☐
Easy to understand	☐	☐	☐	☐	☐
Well organized	☐	☐	☐	☐	☐
Applicable to your tasks	☐	☐	☐	☐	☐

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
SA22-7853-02

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

SA22-7853-02

Cut or Fold
Along Line

Program Number: 5765–E88

SA22-7853-02