Sun HPC ClusterTools™ 5 Software Installation Guide

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054 U.S.A.
650-960-1300

Part No. 817-0082-10
February 2003, Revision A

Send comments about this document to: docfeedback@sun.com

Contents

Preface

This document describes the procedure for installing Sun HPC ClusterTools™ 5 software.

These instructions are intended for an experienced system administrator. For example, to install the cluster tools software on one or more nodes, you must be familiar with the following tasks in a Solaris™ 8 or 9 operating environment:

- Logging in as superuser
- Using the `df` command to check disk space
- Mounting a CD-ROM (using `volcheck` or `mount`)
- Starting and stopping daemons using entries in `/etc/init.d`
- Reading `/var/adm/messages` for possible error messages and debugging information
- Exporting and mounting an NFS file system and using commands and scripts, such as `mount`, `share`, `/etc/init.d/nfs.server`
- Enabling superuser login access to a server
- Setting directory and file permissions to allow read and write access

Before You Read This Book

To follow the procedures described in this document, you should be familiar with the related topics discussed in the following documents:

- The *Sun HPC ClusterTools 5 Software Release Notes*
- Documentation that accompanied your Sun Enterprise™ server or other Sun UltraSPARC™-based server
- Documentation for the Solaris operating environment

- If you are using one of the Distributed Resource Managers: SGE, LSF, or PBS, the associated documentation

How This Book Is Organized

Chapter 1 provides an overview of the Sun HPC ClusterTools software installation utilities.

Chapter 2 describes various installation dependencies.

Chapter 3 explains how to use the graphical installation wizard `ctgui` to install, activate, deactivate, and remove the Sun HPC ClusterTools software.

Chapter 4 explains how to use the Sun HPC ClusterTools software command-line interface (CLI) tools to install, activate, deactivate, and remove the software. It also explains how to use the CLI tools to start and stop all Sun HPC ClusterTools software daemons.

Chapter 5 explains how to use the CLI to install and set up the Sun HPC ClusterTools software on an NFS server.

Chapter 6 describes a set of post-installation procedures you can perform to verify basic functionality of the cluster.

Appendix A summarizes file and directory settings.

Appendix B lists the Sun HPC ClusterTools software installation packages and their sizes.

Using UNIX Commands

This document might not contain information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices.

See one or more of the following for this information:

- *Solaris Handbook for Sun Peripherals*
- AnswerBook2™ online documentation for the Solaris operating environment
- Other software documentation that you received with your system

Typographic Conventions

Typeface*	Meaning	Examples
`AaBbCc123`	The names of commands, files, and directories; on-screen computer output	Edit your `.login` file. Use `ls -a` to list all files. `% You have mail.`
`AaBbCc123`	What you type, when contrasted with on-screen computer output	`% su` `Password:`
AaBbCc123	Book titles, new words or terms, words to be emphasized. Replace command-line variables with real names or values.	Read Chapter 6 in the *User's Guide.* These are called *class* options. You *must* be superuser to do this. To delete a file, type `rm` *filename.*

* The settings on your browser might differ from these settings.

Shell Prompts

Shell	Prompt
C shell	*machine-name*%
C shell superuser	*machine-name*#
Bourne shell and Korn shell	$
Bourne shell and Korn shell superuser	#

.

Related Documentation

Application	Title	Part Number
Sun HPC ClusterTools software documentation	*Read Me First: Guide to Sun HPC ClusterTools Software Documentation*	817-0096-10
Sun HPC ClusterTools software	*Sun HPC ClusterTools 5 Software Release Notes*	817-0081-10
	Sun HPC ClusterTools 5 Software Performance Guide	817-0090-10
	Sun HPC ClusterTools 5 Software Administrator's Guide	817-0083-10
	Sun HPC ClusterTools 5 Software User's Guide	817-0084-10
Sun MPI software programming	*Sun MPI 6.0 Software Programming and Reference Guide*	817-0085-10
Sun S3L software	*Sun S3L 4.0 Software Programming Guide*	817-0086-10
	Sun S3L 4.0 Software Reference Manual	817-0087-10
Prism™ graphical programming environment	*Prism 7.0 Software User's Guide*	817-0088-10
	Prism 7.0 Software Reference Manual	817-0089-10

Accessing Sun Documentation

You can view, print, or purchase a broad selection of Sun documentation, including localized versions, at:

`http://www.sun.com/documentation`

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can email your comments to Sun at:

docfeedback@sun.com

Please include the part number (817-0082-10) of your document in the subject line of your email.

CHAPTER 1

Introduction

Overview

This manual explains how to use the Sun HPC ClusterTools software installation utilities to install, activate, deactivate, and remove Sun HPC ClusterTools software on one or more cluster nodes. You can access these utilities through either a graphical user interface (GUI) or a set of command-line interface (CLI) tools. See TABLE 1-1 for a summary of these interfaces.

TABLE 1-1 Sun HPC ClusterTools Software Installation Utilities

Interface	Role
GUI Utility	
`ctgui`	Graphical interface to a set of Java-based wizards, which are used to install, remove, activate, and deactivate the software on cluster nodes.
CLI Utilities	
`ctinstall`	Installs software on cluster nodes.
`ctremove`	Removes software from cluster nodes.
`ctact`	Activates software.
`ctdeact`	Deactivates software.
`ctstartd`	Starts all Sun HPC ClusterTools software daemons.
`ctstopd`	Stops all Sun HPC ClusterTools software daemons.
`ctnfssvr`	Sets up Sun HPC ClusterTools software on an NFS server.

Note – The Sun HPC ClusterTools 5 software installation utilities offer the same capabilities as the Sun HPC ClusterTools 4 software installation tools, including support for NFS client-server configurations and initiating parallel installation on multiple nodes from a central host. However, the version 5 methods differ from version 4 in several important ways.

Sun HPC ClusterTools 5 Software Packaging Note

Sun HPC ClusterTools 5 software packages are organized into two categories to facilitate installing the software in an NFS client/server configuration. The categories are:

- Root packages – These packages contain daemon startup scripts, which are placed in `/etc/opt` of a cluster node. Each cluster node must have the root packages installed locally.
- Nonroot packages – These packages contain all other Sun HPC ClusterTools 5 software.

The next section explains how the root and non-root packages are handled in NFS and Non-NFS cluster configurations.

NFS and Non-NFS Configurations

You can install Sun HPC ClusterTools software in an NFS client/server configuration or locally on the cluster nodes.

NFS Configuration

In an NFS-based cluster, the tools install the root packages locally on each cluster node using either `ctgui` or `ctinstall` and the non-root packages on the NFS server using the CLI's `ctnfssvr` command. The non-root packages are remotely mounted on the NFS client nodes in the cluster.

The NFS server can, but need not be one of the NFS client nodes in the cluster. When a Sun HPC ClusterTools software NFS server is also a cluster node, both root and non-root packages are installed on the NFS server by means of their respective installation tools. In other words, the non-root packages are installed on the NFS server with `ctnfssvr` and the root packages are installed with either `ctgui` or `ctinstall -c`.

Non-NFS Configuration

In a non-NFS cluster configuration, the tools install a complete copy of the Sun HPC ClusterTools software, both root and non-root packages, locally on each node in the cluster.

Centralized and Local Command Initiation

You can choose between two methods of initiating operations on the cluster nodes:

- Local – Initiate commands directly on the node you are logged in to. The effects of the command are restricted to the local node.
- Centralized – Initiate commands from a central host, specifying the nodes on which the command is to take effect. The initiating host establishes remote connections to the target nodes and broadcasts the commands to them over an rsh, ssh, or telnet connection. The central (initiating) host can be part of the cluster or it can be an administrative system external to the cluster.

Note – Centralized operations are performed on the specified nodes in parallel. That is, when a command is specified on the central host, the operation is initiated on all the specified nodes at the same time.

Support for centralized command initiation is built into the Sun HPC ClusterTools software installation utilities. Issuing these commands from a central host has the equivalent effect as invoking the commands locally using one of the Cluster Console tools, `cconsole`, `ctelnet`, or `crlogin`.

Cluster Console is part of the Sun Cluster software distribution. It is also included in the Sun HPC ClusterTools software distribution as an auxiliary convenience.

Note – Although you do not need Cluster Console for centralized installation of Sun HPC ClusterTools software, it can be a convenient tool for various administrative operations, such as booting cluster nodes in a Custom JumpStart installation or modifying the cluster configuration file, `hpc.conf`.

Centralized command initiation provides a command execution summary for each initiated activity. This summary reports the results of the operation, identifying the nodes on which the operation was successful and the nodes on which it failed. Such reports are not available when using the Cluster Console tools.

The Sun HPC ClusterTools software CLI utilities provide several options that are specific to the centralized command initiation mode and are intended to simplify management of parallel installation of the software from a central host. These options support:

- Creating corresponding versions of local log files on the central host for easier access
- Generating a list of nodes that had successful operations and another list of nodes that were unsuccessful. These pass/fail node lists can then be used in subsequent operations, such as software activation, deactivation, and removal.

The initiating system can be one of the cluster nodes or it can be external to the cluster. It must be a Sun system running the Solaris 8 or 9 operating environment.

Support for Alternative Sun HPC ClusterTools Software Versions

The Sun HPC ClusterTools 5 software can be installed on systems that have Sun HPC ClusterTools 4 software installed and active. This is referred to as *alternative version* support. It is equivalent to the Sun HPC ClusterTools 4 software feature known as *coexistence*, which allowed Sun HPC ClusterTools 4 software to share space on cluster nodes with Sun HPC ClusterTools 3.1 software.

As with coexistence, alternate versions support means that both Sun HPC ClusterTools software versions 5 and 4 can reside on the same cluster nodes, but only one version can be active (available for use) at a given time.

Note – You must use `install_gui` or the other installation tools supplied with the Sun HPC ClusterTools 4 software distribution to perform operations on Sun HPC ClusterTools 4 software. These tools are available on the Sun HPC ClusterTools 4 software distribution CD or in `/opt/SUNWhpc/HPC4.0/bin/Install_Utilities/bin` on an installed system.

Support for Other Installation Contexts

The Sun HPC ClusterTools software installation utilities are completely self-contained and are fully capable of scaling from single node installations to installing very large clusters. However, if you customarily use Custom JumpStart™ and/or Solaris™ Web Start Flash methods for installing software on your servers, you can easily integrate the CLI installation tool `ctinstall` into those contexts.

The following variations on a basic Sun HPC ClusterTools software installation are described briefly below. More detailed descriptions are provided later in the manual.

- Custom JumpStart installation
- Custom JumpStart installation, using Web Start Flash archives

Custom JumpStart Installation

For this context, you must set up a Custom JumpStart environment in advance of the installation. Then you invoke local installation of Sun HPC ClusterTools software on the cluster nodes by integrating the `ctinstall` command in the Custom JumpStart finish script, using the `-l` and `-R` switches.

Custom JumpStart installations are initiated from the console of each Custom JumpStart client. The `cconsole` tool that is included in the Cluster Console software allows you to access multiple consoles through a single common window.

Web Start Flash Installations

The first step for this type of installation is to perform a local installation of the Sun HPC ClusterTools software on a node that will serve as the Flash master. You can use either `ctgui` or `ctinstall -l` for this step. Once the Flash master is fully installed and activated, you create a flash archive and apply it to the target nodes, usually in a Custom JumpStart environment.

This flash archive-based approach creates clones of the Flash master, which includes reinstalling the Solaris operating environment on each clone.

Note – Web Start Flash installations are restricted to cluster environments where all the systems have identical hardware and software configurations.

Sun HPC ClusterTools 5 Installation Log Files

The Sun HPC ClusterTools 5 installation tools log information about installation-related tasks locally on the nodes where installation tasks are performed. The default location for the log files is `/var/sadm/system/logs/hpc`. If installation tasks are initiated from a central host, a summary log file is also created on the central host.

Local, Node-Specific Log Files

Two types of log files are created locally on each cluster node where installation operations take place.

- Task-specific logs – Separate log files are created for each installation-related task. They are: `ct_install.log`, `ct_remove.log`, `ct_act.log`, `ct_deact.log`, `ct_startd.log`, `ct_stopd.log`, and `ct_nfssvr.log`. These log files contain detailed logging information for the most recent associated task. Each time a task is repeated, its log file is overwritten.
- History log – A `ct_history.log` file is created to store all installation-related tasks performed on the local node. This provides a convenient record of the Sun HPC ClusterTools 5 software installation history on the local node. Each time a new installation task is performed on the node, a new log entry is appended to the history log.

These node specific installation log files are created regardless of the installation method used, local or centralized.

Central Node Summary Log

When installation tasks are initiated from a central host, a summary log file named `ct_summary.log` is created on the central host. This log file records the final summary report that is generated by the CLI and GUI. The `ct_summary.log` is not overwritten when a new task is performed. As with the `ct_history.log` file, new entries are appended to the summary log file.

Summary of Tasks Performed by Sun HPC ClusterTools Software Utilities

FIGURE 1-1 shows an overview of the installation-related tasks you can perform using either the graphical installation tool, `ctgui`, or the corresponding command-line interface (CLI) tools. The CLI also supports starting and stopping all Sun HPC ClusterTools software daemons, two operations that are not available through `ctgui`.

Both `ctgui` and the CLI tools require superuser privileges to execute.

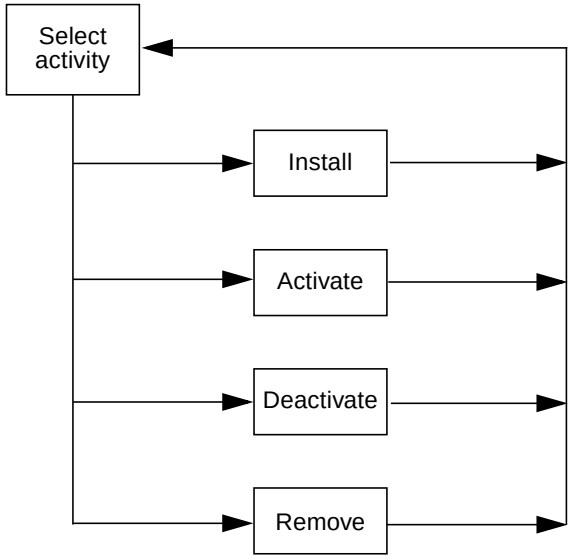

FIGURE 1-1 Sun HPC ClusterTools Software Installation Tasks

Select an Activity

Whether you use `ctgui` or the CLI tools, the various installation-related operations are independent of each other. With `ctgui`, you are given a wizard screen that asks you to select one of the four activities shown in FIGURE 1-1. With the CLI, you simply start the applicable utility: `ctinstall`, `ctact`, `ctdeact`, or `ctremove`. The CLI also provides a pair of utilities, `ctstartd` and `ctstopd`, for starting and stopping all Sun HPC ClusterTools software daemons on the cluster nodes. The GUI does not offer an equivalent feature.

The operations these tools control are described below.

Install the Software

The installation activity loads the Sun HPC ClusterTools software onto cluster nodes. You have the following choices in how the installation will be carried out:

- Install in an NFS client-server cluster configuration.
- Install in a non-NFS cluster configuration.

- Activate the software on the cluster nodes automatically when the installation completes.
- Activate the software on the cluster nodes manually a later time.

These choices are described more fully below.

If HPC ClusterTools 4 software is present on the cluster nodes, you can install the Sun HPC ClusterTools 5 software even if the version 4 software is active. However, version 4 must be inactive before you can activate version 5. Use the version 4 installation/configuration tools provided with the Sun HPC ClusterTools 4 software to change the version 4 software.

With the CLI command `ctinstall`, you can install individual Sun HPC ClusterTools 5 software packages instead of the entire software suite. This capability is not available with `ctgui`.

NFS Installation

In an NFS installation, all Sun HPC ClusterTools software packages, except for a subset containing Sun HPC ClusterTools software daemon startup scripts, are installed on a Sun NFS server and remotely mounted on the NFS client nodes in the cluster. The packages containing the daemon startup scripts must be installed locally on each cluster node.

The packages that contain the daemon startup scripts are placed in the root file system of the cluster nodes and are referred to as the *root* packages.

The NFS server can be one of the NFS client nodes in the cluster, but need not be. When a Sun HPC ClusterTools software NFS server is also one of the cluster nodes, both root and non-root packages are installed on the server.

Non-NFS Installation

In non-NFS configurations, a complete copy of the Sun HPC ClusterTools software—both root and non-root packages—is installed locally on each node in the cluster.

Install With Automatic Activation

When you initiate a software installation operation you can specify to have the nodes activated automatically as soon as the installation process completes.

When you use the automatic activation feature, you must also specify one of the cluster nodes to serve as master node for the cluster. This node will have special daemons running on it. It will also contain cluster configuration and security information not present on the other cluster nodes.

Install Without Activating

You can install the software without activating it. In such cases, you do not specify a master node until you activate the node manually at a later time.

Activate the Sun HPC ClusterTools Software

The node activation step prepares the cluster nodes so that the installed Sun HPC ClusterTools software can be used. Specifically, node activation:

- Sets up symbolic links that point to the Sun HPC ClusterTools software.
- Starts a set of Sun HPC ClusterTools software daemons.

If a node contains both versions 4 and 5 of the Sun HPC ClusterTools software, only one version can be active at a time.

Deactivate the Sun HPC ClusterTools Software

Node deactivation stops the Sun Cluster Runtime Environment (CRE) daemons and removes the symbolic links that point to the Sun HPC ClusterTools software. This step makes the software unusable until the node is activated again.

Remove the Sun HPC ClusterTools Software

This operation deletes Sun HPC ClusterTools software packages from the cluster nodes on which it is executed. If a node is active at the time you initiate the removal operation, it will be deactivated automatically before the software is removed.

With the CLI command `ctremove`, you can remove individual Sun HPC ClusterTools software packages instead of the entire software suite.

Start and Stopping Sun HPC ClusterTools Software Daemons

Starting and stopping the Sun HPC ClusterTools software daemons are subtasks in software activation and deactivation. You can use the commands `ctstartd` and `ctstopd` to start and stop all the Sun HPC ClusterTools software daemons without performing the full range of actions associated with activating and deactivating the software.

Recommendations for Installing Sun HPC ClusterTools 5 Software on Large Cluster Installations

The following are tips for installing Sun HPC ClusterTools 5 software on clusters containing hundreds of nodes using the centralized method:

- Minimize other system activity during installation – Invoking installation of Sun HPC ClusterTools 5 software on hundreds of nodes from a central host imposes high demands on system resources. Avoid system resource exhaustion by keeping the cluster nodes as quiescent as possible during the installation.
- Use a node list file – For various centralized installation tasks, you specify the nodes on which the task is to be invoked. You have the choice of specifying the nodes either on the command line, using the `-n` option or by referencing a node list file using the `-N` option. If you reference a node list file, you only enter the node names once when you create the file.
- Reduce system resource consumption on the central host – You can avoid overtaxing system resources on a single central host by using more than one central host. Simply divide the total list of nodes to be installed on into separate node lists, and initiate the installation commands on the various central hosts, with each host using a different node list.
- Use the `-g` option with CLI-initiated tasks – Use the `-g` option with CLI commands to obtain a list of nodes that successfully executed the command and a separate list of nodes that failed. You can then reference the list of failed nodes with the `-N` option in a later retry of the command.
- Use the `-k` option with CLI-initiated tasks – Use the `-k` option with CLI commands to have all logs saved on the central node where the command was initiated. This option makes it unnecessary to go to each node to examine local logs. The Save Logs feature on the GUI provides the same capability.

Note – When using the telnet connection method, all nodes specified in a given centralized invocation must share the same password. If you have more than one password for the nodes in the cluster, separate the nodes into installation groups in which all nodes share the same password, and perform separate command invocations for each node group.

CHAPTER 2

Preparing for Installation

Before installing Sun HPC ClusterTools software, you need to ensure that the hardware and software that make up your cluster meet certain requirements. These requirements are described in the following sections of this chapter:

General Prerequisites

Before installing Sun HPC ClusterTools software, be sure that all hardware is installed and configured. The system from which the installation will be initiated must have TCP/IP network connections to the target nodes. If the cluster will operate in an NFS configuration, the NFS server must also be connected to the cluster nodes.

TABLE 2-1 outlines additional prerequisites for installation.

TABLE 2-1 General Prerequisites for Installing Sun HPC ClusterTools Software

Description	Requirement
Hardware	Sun UltraSPARC based systems.
Disk space*	Approximately 92 Mbytes per node.
Operating environment	Solaris 8 update 7 or any subsequent Solaris release that supports Sun HPC ClusterTools 5 software.
Java Runtime Environment†	JRE 1.2.0 or above is required for using the installation tool's graphic user interface, `ctgui`. It is not included in Sun HPC ClusterTools 5 software.

TABLE 2-1 General Prerequisites for Installing Sun HPC ClusterTools Software *(Continued)*

Description	Requirement
Resource management	Sun HPC ClusterTools 5 software includes Sun CRE for launching and control of parallel jobs. Sun CRE also provides open integration facilities with the Distributed Resource Management (DRM) systems: Sun Grid Engine (SGE), Load Sharing Facility (LSF), and OpenPBS or Portable Batch System (PBS).
Sun MPI dependence on `libtnfprobe.so`	Sun MPI is linked against `libtnfprobe.so`. This means the packages `SUNWtnfc` and `SUNWtnfcx` must be installed on the system for 32- and 64-bit versions of the library. These packages are included in the Solaris 8 and Solaris 9 software distributions and are installed with the Entire Distribution or Developer System Support installation choices.

`/etc/init.d/nfs.server` file settings	For installations from an NFS server, increase the number of concurrent server requests allowed from 16 to 256. To do this, change the relevant line in `/etc/init.d/nfs.server` (Solaris 8 distribution) or `/etc/default/nfs` (Solaris 9 distribution) and restart the NFS server:
	
	`if [$startnfsd -ne 0];`
	`then`
	`  /usr/lib/nfs/mountd`
	`  /usr/lib/nfs/nfsd -a 256`
	`fi`

`/etc/system` file settings	Edit the `/etc/system` file on every node in your cluster. Add the following entry to the file on each node and then reboot the nodes: `set pt_cnt=1024`
Coexists with Sun HPC ClusterTools 4 software only	Remove any version of Sun HPC ClusterTools software earlier than version 4.

* If the Sun HPC ClusterTools software is installed in an NFS-based cluster configuration, the packages installed on each NFS client node require only 146 kilobytes. If the NFS server is not also a cluster node, the Sun HPC ClusterTools software packages installed on it require 91957 kilobytes. If the NFS server is also a cluster node, it requires the full 92,103 kilobytes.

† If you run `ctgui` in a GNOME desktop environment, the `ctgui` window may require resizing by hand after it opens. See "Dependencies on GNOME and JDK Versions" on page 15 for more information.

The Sun HPC ClusterTools software installation process sets certain parameters in `/etc/system` to enable RSM communication. If you wish to change the default settings to meet other requirements of your site, please see the *Sun HPC ClusterTools Software Administrator's Guide* for detailed information.

Sun S3L Requirement

The Sun Scalable Scientific Subroutine Library (Sun S3L) uses the Sun Performance Library for node-level operations. This means either Forte™ Developer 6, update 2 or Sun One Studio 7, Compiler Collection must be present.

Supported Compilers

Sun HPC ClusterTools software supports Forte Developer 6, update 2 and Sun One Studio 7 Compiler Collection for C, C++, and Fortran compilers.

Cluster Console

The Cluster Console tools, `cconsole`, `ctelnet`, and `crlogin`, are available on the Sun HPC ClusterTools software distribution in the package `SUNWccon`. Two versions of this package are provided in `/cdrom/hpc/Cluster_Console/Solaris_8` and `/cdrom/hpc/Cluster_Console/Solaris_9`.

Use `pkgadd` to install the appropriate version of `SUNWccon` from the Sun HPC ClusterTools 5 software CD-ROM. You cannot install `SUNWccon` using `ctgui` or `ctinstall`.

Dependencies on GNOME and JDK Versions

If you plan to use `ctgui` in a GNOME desktop environment, be certain to use JDK version 1.4.0. This will avoid a sizing problem that prevents the wizard from opening to its correct size. This problem is fixed in JDK 1.4.0.

If you must use `ctgui` in a GNOME desktop environment with a version of JDK earlier than 1.4.0, you can resize the window to the desired dimensions by hand.

CHAPTER 3

Using ctgui

This chapter explains how to use the graphical interface tool, `ctgui`, to install, activate, deactivate, and remove Sun HPC ClusterTools software on cluster nodes.

If you plan to install the software in an NFS cluster configuration, keep the following in mind:

- You will use the CLI command `ctnfssvr` to set up the Sun HPC ClusterTools software on the NFS server and install the non-root software packages on it.
- If you plan to use the NFS server as a cluster node, use `ctgui` or `ctinstall -c` to install the root packages on it.

See Chapter 5 for instructions on setting up and installing software on NFS servers.

Note – If you use rsh connections for centralized operations on hundreds of nodes at a time, the operations may encounter system resource limitations that prevent the connections from being established to all the nodes. For clusters with hundreds of nodes, it is best to perform these operations on subsets of nodes, one subset at a time, with no more than 200 nodes in a subset.

At the end of each operation, a panel will be displayed that summarizes the results of the operation on all the nodes. You can also view and save to a central location the log files from each node that participated in the operation.

▼ To Start `ctgui`

1. **Boot the cluster nodes and load a CD-ROM containing the Sun HPC ClusterTools software in each. For centralized installations, do this on the central host as well.**
2. **Log in as superuser on the system from which you will be initiating the command—the central host or directly on a cluster node if initiating the command locally.**

3. **If the Sun HPC ClusterTools software has *not* been installed yet, change directory to** *distribution*`/hpc/Product/Install_Utilities/bin`**, where** *distribution* **is the location of the Sun HPC ClusterTools software CD-ROM. Otherwise, go to Step 4.**

4. **If the software was previously installed and you intend to perform such tasks as activation, deactivation, or removal of the software, change directory to** `$INSTALL_LOC/SUNWhpc/HPC5.0/bin/Install_Utilities/bin`**, where** `$INSTALL_LOC` **is the location where the software was installed.**

5. **Start** `ctgui`**, which causes the first GUI panel, labeled Overview, to be displayed.**

```
# ./ctgui
```

Note – The `ctgui` is fully self-documented so you may not need to read the following procedure. However, it is provided as an optional reference.

▼ To Install Sun HPC ClusterTools Software Using `ctgui`

This procedure assumes that `ctgui` has been started and that you are looking at the first GUI panel.

1. **Read the Overview text and click on the Next button to display the Select an Activity panel.**

   FIGURE 3-1 illustrates the sequence of `ctgui` panels for the Install activity, beginning with the Select an Activity panel.

2. **Select Install and click on the Next button.**

   This begins a sequence of installation-related GUI panels, some of which request information about the installation to be performed.

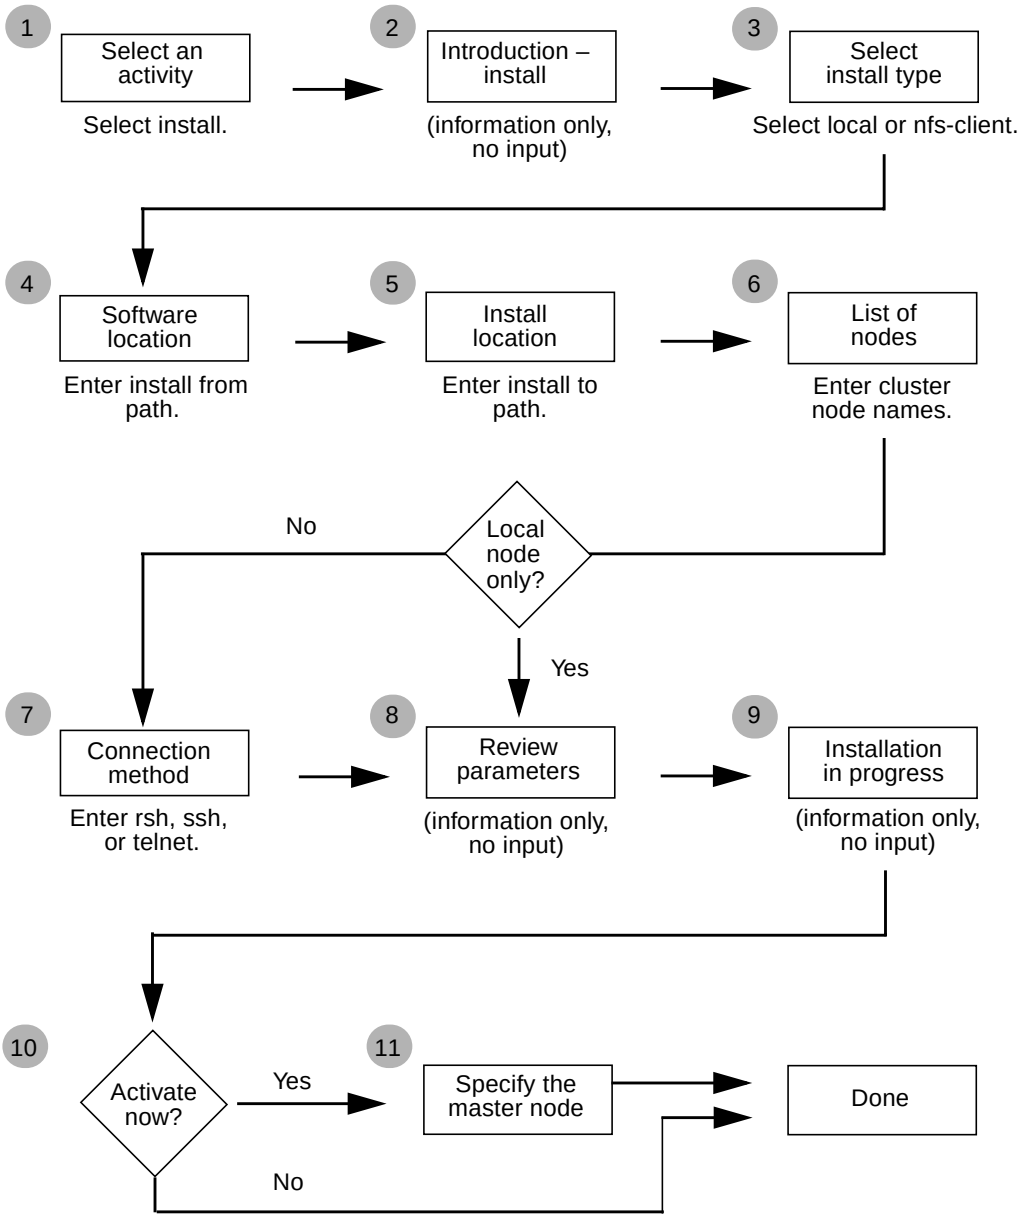

FIGURE 3-1 Sequence of `ctgui` Panels for Installing Sun HPC ClusterTools Software

3. **Select the installation type – either local or nfs-client – and click on the Next button.**

   If you select local, the full set of Sun HPC ClusterTools software packages will be installed. If you select nfs-client, only the root packages will be installed.

4. **Specify the location of the software to be installed. To install from the default location, simply click on the Next button.**

5. **Specify the location where you want the software to reside. To use the default location, `/opt`, click on the Next button.**

6. **Specify the cluster nodes on which the installation is to take place, using one of the following methods:**

   - Type a list of nodes, separating the names with commas, and click on Add.
   - To specify the node you are logged in to, click the Add Local Node button.
   - To load a file containing the list of target nodes, click the Load File ... button. Then specify the full path of the node list file.
   - You can delete nodes from the resulting node list with Remove and Clear All.

   When you are satisfied with the list of target nodes, click the Next button.

7. **If you specified only the local node in Step 6, skip this step. Otherwise, specify one of the following connection methods:**

   - rsh – Requires a trusted hosts setup.
   - ssh – The initiating node must be able to log in as superuser to the target nodes without being prompted for any interaction, such as a password.
   - telnet – All cluster nodes must share the same password. If some nodes do not use the same password as other nodes, install the software on groups of nodes that use a common password.

8. **Review the values you have specified for the installation.**

   If you want to change any values, use the Back button to step through the previous panels until you are at the panel where you specified the value and make the desired change. Then use the Next button to return by steps to this panel.

   When you are satisfied with the parameter values, click the Finish button to start the installation.

9. **An Installation in Progress... panel will be displayed while the software is being installed.**

   Several forms of status information are provided on this panel:

   - Across the top you are told how many installations have passed, how many have failed, and how many are busy.
   - A separate progress bar and percent complete indication are provided for each node.

- Buttons are provided that allow you to stop the installation on all nodes in the supplied node list or on specific nodes.

You can also view the installation logs for each of the nodes.

When you are satisfied with the installation status, click the Next button, which causes the Is It Time to Activate panel to be displayed.

10. **Click the Activate Now button to make the software ready for immediate use or click the Activate Later button to postpone the activation step to another time.**

11. **If you choose Activate Now, you will be asked to name the master node for the cluster.**

    This is one node in the cluster where the Sun CRE master daemons will run.

    This completes the installation process on the cluster nodes. The next possibilities are:

    - Return to the Select an Activity panel – Click the Next button.
    - Return to the Installation in Progress panel – Click the Back button.
    - Leave the `ctgui` environment – Click the Exit button.

    If you are installing in an NFS cluster configuration and you have not yet set up and installed on the NFS server, go to Chapter 5 for instructions.

▼ To Activate Sun HPC ClusterTools Software Using `ctgui`

The procedure described below assumes that `ctgui` has been started and that you are looking at the first GUI panel.

Note – If Sun HPC ClusterTools 4 software is resident on the cluster nodes, it must be inactive before you perform the following procedure.

Special Note for Using `sunhpc_rhosts`

If you are using the `sunhpc_rhosts` authentication method (the default), all the nodes to be activated must be listed in `/etc/sunhpc_rhosts` on the master node. To ensure that this file is created and automatically maintained by the installation tools, be certain to initiate the activation operation from the master node. If a node is not included in this file on the master node, it will not be activated.

Special Note for Activating Sun HPC ClusterTools Software in an NFS Configuration

When you want to activate Sun HPC ClusterTools software in an NFS configuration, you must ensure that the activation tool is able to locate the software mount point. You can do this in either of the following ways:

- Use `ctgui` under the NFS mount point on an NFS client:
  mount_point`/*/SUNWhpc/HPC5.0/bin/Install_Utilities/bin`
- Or, explicitly mount the directory in which `SUNWhpc` is installed. For example, if `SUNWhpc` is installed in `/export/apps`, enter the following on an NFS client:

```
# mount server:/export/apps mount_point
```

1. **Read the Overview text on the first GUI panel and click on the Next button to display the Select an Activity panel.**
2. **Select Activate and click the Next button.**

   For this task, you will be asked to name a master node for the cluster and the list of nodes that are to be activated. If you are activating nodes from a central host, you will also be asked to specify the connection method.
3. **Specify the master node.**

   If a master node was specified previously, enter that node's name.
4. **Specify the cluster nodes that are to be activated.**
5. **If you specified only the local node in Step 4, skip this step. Otherwise, specify one of the following connection methods:**
   - rsh – Requires a trusted hosts setup.
   - ssh – The initiating node must be able to log in as superuser to the target nodes without being prompted for any interaction, such as a password.
   - telnet – All cluster nodes must share the same password. If some nodes do not use the same password as other nodes, activate the nodes in groups that use a common password.
6. **Review the values you have specified for the software activation.**

   If you want to change any values, use the Back button to step through the previous panels until you are at the panel where you specified the value and make the desired change. Then use the Next button to return by steps to this panel.

   When you are satisfied with the activation parameter values, click the Finish button to start the activation.

7. **An Activation in Progress... panel will be displayed.**

   The following status information is provided by this panel:

   - The number of nodes where activation has successfully completed, the number that have failed, and the number that are busy.
   - A separate progress bar and percent complete indication are provided for each node.
   - Buttons are provided that allow you to stop the activation on all nodes in the supplied node list or on specific nodes.

   You can also view activation logs for each of the nodes.

   When you are satisfied with the activation status, click the Next button, which causes Summary panel to be displayed. At this point, you can:

   - Return to the Select an Activity panel – Click the Next button.
   - Return to the Activation in Progress panel – Click the Back button.
   - Leave the `ctgui` environment – Click the Exit button.

▼ To Deactivate Sun HPC ClusterTools Software Using `ctgui`

This procedure assumes that `ctgui` has been started and that you are looking at the first GUI panel.

1. **Read the Overview text and click on the Next button to display the Select an Activity panel.**
2. **Select Deactivate and click the Next button.**

   For this task, you will be asked for the list of nodes that are to be deactivated. If you are deactivating from a central host, you will also be asked to specify the connection method.

3. **Specify the cluster nodes that are to be deactivated.**
4. **If you specified only the local node in Step 3, skip this step. Otherwise, specify one of the following connection methods**:

   - rsh – Requires a trusted hosts setup.
   - ssh – The initiating node must be able to log in as superuser to the target nodes without being prompted for any interaction, such as a password.
   - telnet – All cluster nodes must share the same password. If some nodes do not use the same password as other nodes, deactivate the nodes groups that use a common password.

5. **Review the values you have specified for the software deactivation.**

   If you want to change any values, use the Back button to step through the previous panels until you are at the panel where you specified the value and make the desired change. Then use the Next button to return by steps to this panel.

   When you are satisfied with the deactivation parameter values, click the Finish button to start the deactivation.

6. **A Deactivation in Progress... panel will be displayed.**

   The following status information is provided by this panel:

   - The number of nodes where deactivation has successfully completed, the number that have failed, and the number that are busy.
   - A separate progress bar and percent complete indication are provided for each node.
   - Buttons are provided that allow you to stop the deactivation on all nodes in the supplied node list or on specific nodes.

   You can also view deactivation logs for each of the nodes.

   When you are satisfied with the deactivation status, click the Next button, which causes Summary panel to be displayed. At this point, you can:

   - Return to the Select an Activity panel – Click the Next button.
   - Return to the Deactivation in Progress panel – Click the Back button.
   - Leave the `ctgui` environment – Click the Exit button.

▼ To Remove Sun HPC ClusterTools Software Using `ctgui`

This procedure assumes that `ctgui` has been started and that you are looking at the first GUI panel.

1. **Read the Overview text and click on the Next button to display the Select an Activity panel.**

2. **Select Remove and click the Next button.**

   For this task, you will be asked for the list of nodes that are to be removed. If you are removing nodes from a central host, you will also be asked to specify the connection method.

3. **Specify the cluster nodes that are to be removed.**

4. **If you specified only the local node in Step 3, skip this step. Otherwise, specify one of the following connection methods:**

   - rsh – Requires a trusted hosts setup.

- ssh – The initiating node must be able to log in as superuser to the target nodes without being prompted for any interaction, such as a password.
- telnet – All cluster nodes must share the same password. If some nodes do not use the same password as other nodes, remove the nodes in groups that use a common password.

5. Review the values you have specified for the software removal.

If you want to change any values, use the Back button to step through the previous panels until you are at the panel where you specified the value and make the desired change. Then use the Next button to return by steps to this panel.

When you are satisfied with the removal parameter values, click the Finish button to start the removal.

6. A Software Removal in Progress... panel will be displayed.

The following status information is provided by this panel:

- The number of nodes where removal has successfully completed, the number that have failed, and the number that are busy.
- A separate progress bar and percent complete indication are provided for each node.
- Buttons are provided that allow you to stop the removal of all nodes in the supplied node list or of specific nodes.

You can also view node removal logs for each of the nodes.

When you are satisfied with the node removal status, click the Next button, which causes Summary panel to be displayed. At this point, you can:

- Return to the Select an Activity panel – Click the Next button.
- Return to the Removal in Progress panel – Click the Back button.
- Leave the `ctgui` environment – Click the Exit button.

CHAPTER 4

Using the CLI Tools On Sun HPC Cluster Nodes

This chapter explains how to use the following Sun HPC ClusterTools software installation utilities.

- `ctinstall` – Install the software on the cluster nodes.
- `ctact` – Activate the software on the cluster nodes.
- `ctdeact` – Deactivate the software on the cluster nodes.
- `ctremove` – Remove the software from the cluster nodes.
- `ctstartd` – Start all Sun HPC ClusterTools software daemons on the cluster nodes.
- `ctstopd` – Stop all Sun HPC ClusterTools software daemons on the cluster nodes.

If you are installing the software in an NFS cluster configuration, remember that:

- You will use the CLI command `ctnfssvr` to set up the NFS server and install the non-root software packages on it.
- If you plan to use the NFS server as a cluster node, use `ctgui` or `ctinstall -c` to install the root packages on it.

See Chapter 5 for instructions on setting up and installing software on NFS servers.

Note – If you use rsh connections for centralized operations on hundreds of nodes at a time, the operations may encounter system resource limitations that prevent the connections from being established to all the nodes. For clusters with hundreds of nodes, it is best to perform these operations on subsets of nodes, one subset at a time, with no more than 200 nodes in a subset.

Special Note for Activating Sun HPC ClusterTools Software in an NFS Configuration

When you want to activate Sun HPC ClusterTools software in an NFS configuration, you must ensure that the activation tool is able to locate the software mount point. You can do this in either of the following ways:

- Use the activation tool under the NFS mount point on an NFS client:

```
mount_point/*/SUNWhpc/HPC5.0/bin/Install_Utilities/bin
```

- Or, explicitly mount the directory in which `SUNWhpc` is installed. For example, if `SUNWhpc` is installed in `/export/apps`, enter the following on an NFS client:

```
# mount server:/export/apps mount_point
```

Initial Steps

The steps below are common to any operation in which you would use CLI commands.

▼ To Access the Sun HPC ClusterTools Software CLI

1. **Load a CD-ROM containing the Sun HPC ClusterTools software in each cluster node. If using a central host for command initiation, do this on the central host as well.**
2. **Log in as superuser.**

   If you are using a central command initiation host, do this step on the central host. If operating in direct local mode, log in as superuser on the cluster node.
3. **If the Sun HPC ClusterTools software has *not* already been installed, change directory to** *distribution*`/hpc/Product/Install_Utilities/bin`**, where** *distribution* **is the location of the Sun HPC ClusterTools software CD-ROM. Otherwise, go to Step 4.**

4. **If the software was previously installed and you intend to perform other tasks, such as activation, deactivation, or removal, change directory to** `$INSTALL_LOC/SUNWhpc/HPC5.0/bin/Install_Utilities/bin`**, where** `$INSTALL_LOC` **is the location where the software was installed.**

   You can now start using the CLI commands. They are described separately below, with examples of common applications given for each.

   For usage information on any command, either enter the command without options or with the `-h` option.

```
# ./command
or
# ./command -h
```

Install Sun HPC ClusterTools 5 Software

Use the `ctinstall` command to install Sun HPC ClusterTools software on cluster nodes. See TABLE 4-1 for a summary of the `ctinstall` options. Explanations of their use are provided in the following contexts:

- Centralized operations in non-NFS configurations
- Centralized operations in NFS configurations
- Local operations in NFS and non-NFS configurations

TABLE 4-1 `ctinstall` Options

Options	Description
General	
`-h`	Command help.
`-l`	Execute the command on the local node only.
`-R`	Specify the full path to be used as the root path.
`-x`	Turn on command debug at the specified nodes.
Command Specific	
`-a`	Activate automatically after installation completes. Must use with `-m`.
`-c`	Use when installing on an NFS client. Installs only root packages.
`-d`	Specify a non-default *install from* location. The default is *distribution*`/hpc/Product`, relative to the directory where `ctinstall` is invoked.

TABLE 4-1 `ctinstall` Options *(Continued)*

Options	Description
`–m`	Specify a master node. Use this option with `–a`.
`–p`	List of packages to be installed. Separate names by comma.
`–s`	Specify a security option for Sun CRE. The choices are `sunhpc_rhosts`, `rhosts`, `des`, and `krb5`. The default is `sunhpc_rhosts`.
`–t`	Specify a nondefault *install to* location. The default is `/opt`.
Centralized Operations Only	
`–g`	Generate node lists of successful and unsuccessful installations.
`–k`	Specify a central location for storing log files of all specified nodes.
`–n`	List of nodes targeted for installation. Separate names by comma.
`–N`	File containing list of nodes targeted for installation. One node per line.
`–r`	Remote connection method: rsh, ssh, or telnet.
`–S`	Specify full path to an alternate ssh executable.

Installing Software From a Central Host in Non-NFS Configurations

This section shows examples of software installations in which the `ctinstall` command is initiated from a central host in a non-NFS configuration.

Install Without Activating

CODE EXAMPLE 4-1

```
# ./ctinstall –n node1,node2 –r rsh
```

CODE EXAMPLE 4-1 installs the full Sun HPC ClusterTools software suite (root and non-root packages) on `node1` and `node2` from a central host. The node list is specified on the command line. The remote connection method is rsh. This requires a trusted hosts setup.

The software will not be ready for use when the installation process completes. It must be activated by hand before it can be used.

CODE EXAMPLE 4-2

```
# ./ctinstall –n node1,node2 –r ssh
```

CODE EXAMPLE 4-2 is the same as CODE EXAMPLE 4-1, except that the remote connection method is ssh. This method requires that the initiating node be able to log in as superuser to the target nodes without being prompted for any interaction, such as a password.

CODE EXAMPLE 4-3

```
# ./ctinstall –N /tmp/nodelist –r telnet
```

CODE EXAMPLE 4-3 installs the full Sun HPC ClusterTools software suite (root and non-root packages) on the set of nodes listed in the file `/tmp/nodelist` from a central host. A node list file is particularly useful when you have a large set of nodes or you want to run operations on the same set of nodes repeatedly.

The node list file has the following contents:

```
# Node list for CODE EXAMPLE 4-2

node1
node2
```

The remote connection method is telnet. All cluster nodes must share the same password. If some nodes do not use the same password as others, install the software in groups, each group consisting of nodes that use a common password.

The software will not be ready for use when the installation process completes. It must be activated by hand before it can be used.

CODE EXAMPLE 4-4

```
# ./ctinstall –N /tmp/nodelist –r telnet –k /tmp/cluster-logs –g
```

CODE EXAMPLE 4-4 is the same as CODE EXAMPLE 4-3, except it includes the –k and –g options.

In this example, the –k option causes the local log files of all specified nodes to be saved in `/tmp/cluster-logs` on the central host.

Note – Specify a directory that is local to the central host rather than an NFS-mounted directory. This will avoid unnecessary network traffic in the transfer of log files and will result in faster execution of the operation.

The `-g` option causes a pair of node list files to be created on the central host in `/var/sadm/system/logs/hpc/nodelists`. One file, `ctinstall.pass$$`, contains a list of the nodes on which the installation was successful. The other file, `ctinstall.fail$$`, lists the nodes on which the installation was unsuccessful. The `$$` symbol is replaced by the process number associated with the installation.

These generated node list files can then be used for command retries or in subsequent operations using the `-N` switch.

CODE EXAMPLE 4-5

```
# ./ctinstall -N /tmp/nodelist -r telnet -p SUNWcremn,SUNWmpimn
```

CODE EXAMPLE 4-5 installs the packages `SUNWcremn` and `SUNWmpimn` on the set of nodes listed in the file `/tmp/nodelist`. No other packages are installed. The remote connection method is telnet.

The `-p` option can be useful if individual packages were not installed on the nodes by `ctinstall`.

Caution – Do not use the `-p` option on NFS client nodes—that is, in conjunction with the `-c` option. The `-c` option is responsible for controlling which Sun HPC ClusterTools software packages belong on NFS client nodes.

Install and Activate Automatically

Note – If you are using the `sunhpc_rhosts` authentication method, all the nodes to be activated must be listed in `/etc/sunhpc_rhosts` on the master node. To ensure that this file is created and automatically maintained by the installation tools, be certain to initiate the activation operation from the master node. If a node is not included in this file on the master node, it will not be activated.

CODE EXAMPLE 4-6

```
# ./ctinstall -N /tmp/nodelist -r rsh -a -m node2
```

CODE EXAMPLE 4-6 installs the full Sun HPC ClusterTools software suite (root and non-root packages) on the nodes listed in the file `/tmp/nodelist`. The remote connection method is rsh.

The software will be activated automatically as soon as the installation is complete. Because activation is automatic, a master node must be specified for the cluster in advance. This is `node2` in CODE EXAMPLE 4-6. If a master node is not specified, an error message is displayed.

Installing Software From a Central Host in NFS Configurations

This section shows examples of software installations in which the `ctinstall` command is initiated from a central host in an NFS configuration.

Install Without Activating

CODE EXAMPLE 4-7

```
# ./ctinstall –c –n node1,node2 –r rsh
```

CODE EXAMPLE 4-7 is the same as CODE EXAMPLE 4-1, except that `node1` and `node2` are NFS client nodes. The `–c` option causes only root packages to be installed on these nodes. If the NFS server is to be used as a cluster node, run this command on it as well.

Use `ctnfssvr` to set up the NFS server and install the non-root packages on it.

Install and Activate Automatically

CODE EXAMPLE 4-8

```
# ./ctinstall –c –n node1,node2 –r rsh –a –m node2
```

CODE EXAMPLE 4-8 is the same as CODE EXAMPLE 4-7, except it includes the options `–a` and `–m`, which cause the software to be activated automatically and specify the cluster's master node, respectively.

Note – Since this command will activate the software on NFS client nodes as soon as the installation completes, the NFS server must be properly installed and enabled before this operation is performed. See Chapter 5 for details on NFS server setup operations.

Installing Software Locally in Non-NFS Configurations

This section shows examples of software installations in which the `ctinstall` command is initiated on the local node in non-NFS configurations.

Note – The options `–g`, `–k`, `–n`, `–N`, `–r`, and `–S` are incompatible with local (non-centralized) installations. If the `–l` option is used with any of these options, an error message is displayed.

Install Locally Without Activating

CODE EXAMPLE 4-9

```
# ./ctinstall –l
```

CODE EXAMPLE 4-9 installs the full Sun HPC ClusterTools software suite (root and non-root packages) on the local node only.

CODE EXAMPLE 4-10

```
# ./ctinstall –l –p SUNWcremn,SUNWmpimn
```

CODE EXAMPLE 4-10 installs the packages `SUNWcremn` and `SUNWmpimn` on the local node.

Install Locally and Activate Automatically

CODE EXAMPLE 4-11

```
# ./ctinstall –l –a  –m node2
```

CODE EXAMPLE 4-11 installs the full Sun HPC ClusterTools software suite (root and non-root packages) on the local node and causes it to be activated as soon as the installation is complete. It also specifies the cluster master node as `node2`.

Note – The local node needs to be told which cluster node is the master node.

Installing Software Locally in NFS Configurations

This section shows examples of software installations in which the `ctinstall` command is initiated on the local node in NFS configurations.

Install Locally Without Activating

CODE EXAMPLE 4-12

```
# ./ctinstall –c –l
```

CODE EXAMPLE 4-12 installs the Sun HPC ClusterTools software root packages on the local node.

Install Locally and Activate Automatically

CODE EXAMPLE 4-13

```
# ./ctinstall –c –l –a –m node2
```

CODE EXAMPLE 4-13 is the same as CODE EXAMPLE 4-12, except the software is activated as soon as the installation completes. The NFS server must be installed and enabled before this step can be taken.

Activate Sun HPC ClusterTools Software

Use the `ctact` command to activate Sun HPC ClusterTools software on cluster nodes. See TABLE 4-2 for a summary of the `ctact` options.

Note – The general options and options specific to centralized operations serve essentially the same role for `ctact` as for `ctinstall`. Consequently, fewer examples are used to illustrate `ctact` than were used for `ctinstall`.

TABLE 4-2 `ctact` Options

Options	Description
General	
`-h`	Command help.
`-l`	Execute the command on the local node only.
`-R`	Specify the full path to be used as the root path.
`-x`	Turn on command debug at the specified nodes.
Command Specific	
`-m`	Specify a master node.
Centralized Operations Only	
`-g`	Generate node lists of successful and unsuccessful activation.
`-k`	Specify a central location for storing copies of local log files.
`-n`	List of nodes targeted for activation. Separate names by comma.
`-N`	File containing list of nodes targeted for activation. One node per line.
`-r`	Remote connection method: rsh, ssh, or telnet.
`-S`	Specify full path to an alternate ssh executable.

Note – If you are using the `sunhpc_rhosts` authentication method, all the nodes to be activated must be listed in `/etc/sunhpc_rhosts` on the master node. To ensure that this file is created and automatically maintained by the installation tools, be certain to initiate the activation operation from the master node. If a node is not included in this file on the master node, it will not be activated.

Activating Nodes From a Central Host

This section shows examples of software activation in which the `ctact` command is initiated from a central host.

Activate Specified Cluster Nodes

CODE EXAMPLE 4-14

```
# ./ctact –n node1,node2 –r rsh –m node2
```

CODE EXAMPLE 4-14 activates the software on `node1` and `node2` and specifies `node2` as the master node. The remote connection method is rsh.

CODE EXAMPLE 4-15

```
# ./ctact –n node1,node2 –r rsh –m node2 –k /tmp/cluster-logs –g
```

CODE EXAMPLE 4-15 is the same as CODE EXAMPLE 4-14, except it specifies the options –k and –g.

In this example, the `–k` option causes the local log files of all specified nodes to be saved in `/tmp/cluster-logs` on the central host.

Note – Specify a directory that is local to the central host rather than an NFS-mounted directory. This will avoid unnecessary network traffic and will result in faster execution of the operation.

The `–g` option causes files `ctact.pass$$` and `ctact.fail$$` to be created on the central host in `/var/sadm/system/logs/hpc/nodelists`. `ctact.pass$$` lists the cluster nodes on which software activation was successful and `ctact.fail$$` lists the nodes on which activation was unsuccessful. The `$$` symbol is replaced by the process number associated with the activation.

These generated node list files can then be used for command retries or in subsequent operations using the `–N` switch.

Activating the Local Node

This section shows an example of software activation on the local node.

Activate Locally

CODE EXAMPLE 4-16

```
# ./ctact –l –m node2
```

CODE EXAMPLE 4-16 activates the software on the local node and specifies `node2` as the master node.

Deactivate Sun HPC ClusterTools Software

Use the `ctdeact` command to deactivate Sun HPC ClusterTools software on cluster nodes. See TABLE 4-3 for a summary of the `ctdeact` options.

TABLE 4-3 `ctdeact` Options

Options	Description
General	
–h	Command help.
–l	Execute the command on the local node only.
–R	Specify the full path to be used as the root path.
–x	Turn on command debug at the specified nodes.
Centralized Operations Only	
–g	Generate node lists of successful and unsuccessful deactivation.
–k	Specify a central location for storing copies of local log files.
–n	List of nodes targeted for deactivation. Separate names by comma.
–N	File containing list of nodes to be deactivated. One node per line.
–r	Remote connection method: rsh, ssh, or telnet.
–S	Specify full path to an alternate ssh executable.

Deactivating Software From a Central Host

This section shows examples of software deactivation in which the `ctdeact` command is initiated from a central host.

Deactivate Specified Cluster Nodes

CODE EXAMPLE 4-17

```
# ./ctdeact –N /tmp/nodelist –r rsh
```

CODE EXAMPLE 4-17 deactivates the software on the nodes listed in `/tmp/nodelist`. The remote connection method is rsh.

CODE EXAMPLE 4-18

```
# ./ctact –N /tmp/nodelist –r rsh –k /tmp/cluster-logs –g
```

CODE EXAMPLE 4-18 is the same as CODE EXAMPLE 4-17, except it specifies the options –k and –g.

In this example, the `–k` option causes the local log files of all specified nodes to be saved in `/tmp/cluster-logs` on the central host.

Note – Specify a directory that is local to the central host rather than an NFS-mounted directory. This will avoid unnecessary network traffic in the transfer of log files and will result in faster execution of the operation.

The `–g` option causes files `ctdeact.pass$$` and `ctdeact.fail$$` to be created on the central host. `ctdeact.pass$$` lists the cluster nodes where software deactivation was successful. `ctdeact.fail$$` lists the nodes where deactivation was unsuccessful. The `$$` symbol is replaced by the process number associated with the software deactivation.

These generated node list files can then be used for command retries or in subsequent operations using the `–N` switch.

Deactivating the Local Node

This section shows software deactivation on the local node.

Deactivate Locally

CODE EXAMPLE 4-19

```
# ./ctdeact –l
```

CODE EXAMPLE 4-19 deactivates the software on the local node.

Remove Sun HPC ClusterTools Software

Use the `ctremove` command to remove Sun HPC ClusterTools software from cluster nodes. See TABLE 4-4 for a summary of the `ctremove` options.

Note – If the nodes are active at the time `ctremove` is initiated, they will be deactivated automatically before the removal process begins.

TABLE 4-4 `ctremove` Options

Options	Description
General	
–h	Command help.
–l	Execute the command on the local node only.
–R	Specify the full path to be used as the root path.
–x	Turn on command debug at the specified nodes.
Command Specific	
–p	List of packages to be selectively removed. Separate names by comma.
Centralized Operations Only	
–g	Generate node lists of successful and unsuccessful removals.
–k	Specify a central location for storing copies of local log files.
–n	List of nodes targeted for removal. Separate names by comma.
–N	File containing list of nodes targeted for removal. One node per line.
–r	Remote connection method: rsh, ssh, or telnet.
–S	Specify full path to an alternate ssh executable.

Removing Nodes From a Central Host

This section shows examples of software removal in which the `ctremove` command is initiated from a central host.

Remove Software From Specified Cluster Nodes

CODE EXAMPLE 4-20

```
# ./ctremove –N /tmp/nodelist –r rsh
```

CODE EXAMPLE 4-20 removes the software from the nodes listed in `/tmp/nodelist`. The remote connection method is rsh.

CODE EXAMPLE 4-21

```
# ./ctremove –N /tmp/nodelist –r rsh –k /tmp/cluster-logs –g
```

CODE EXAMPLE 4-21 is the same as CODE EXAMPLE 4-20, except it specifies the options –k and –g.

CODE EXAMPLE 4-22

```
# ./ctremove –N /tmp/nodelist –r rsh –p SUNWcremn,SUNWmpimn
```

CODE EXAMPLE 4-22 removes the packages `SUNWcremn` and `SUNWmpimn` from the nodes listed in `/tmp/nodelist`. The remote connection method is rsh.

Removing Software From the Local Node

This section shows software removal from the local node.

Remove Software Locally

CODE EXAMPLE 4-23

```
# ./ctremove –l
```

CODE EXAMPLE 4-23 deactivates the software on the local node.

CODE EXAMPLE 4-24

```
# ./ctremove –l –p SUNWcremn,SUNWmpimn
```

CODE EXAMPLE 4-24 removes the packages SUNWcremn and SUNWmpimn from the local node.

Start Sun HPC ClusterTools Software Daemons

Use the ctstartd command to start all Sun HPC ClusterTools software daemons on the cluster nodes. Once the Sun HPC ClusterTools 5 software is activated, ctstartd is available in /opt/SUNWhpc/sbin.

See TABLE 4-5 for a summary of the ctstartd options.

TABLE 4-5 ctstartd Options

Options	Description
General	
–h	Command help.
–l	Execute the command on the local node only.
–R	Specify the full path to be used as the root path.
–x	Turn on command debug at the specified nodes.
Centralized Operations Only	
–g	Generate node lists of successful and unsuccessful ctstard operations.
–k	Specify a central location for storing copies of local log files.
–n	List of nodes where daemons will be started. Separate names by comma.
–N	File containing list of nodes where daemons will be started. One node per line.
–r	Remote connection method: rsh, ssh, or telnet.
–S	Specify full path to an alternate ssh executable.

Starting Daemons From a Central Host

This section shows how to start Sun HPC ClusterTools software daemons from a central host.

Start Daemons on Specified Cluster Nodes

CODE EXAMPLE 4-25

```
# ./ctstartd –N /tmp/nodelist –r rsh
```

CODE EXAMPLE 4-25 starts the Sun HPC ClusterTools software daemons on the nodes listed in `/tmp/nodelist`. The remote connection method is rsh.

CODE EXAMPLE 4-26

```
# ./ctstartd –N /tmp/nodelist –r rsh –k /tmp/cluster-logs –g
```

CODE EXAMPLE 4-26 is the same as CODE EXAMPLE 4-25, except it specifies the options `–k` and `–g` to gather log information centrally and to generate pass and fail node lists.

Starting Daemons on the Local Node

Start Daemons Locally

CODE EXAMPLE 4-27

```
# ./ctstartd –l
```

CODE EXAMPLE 4-27 starts the Sun HPC ClusterTools software daemons on the local node.

Stop HPC ClusterTools Software Daemons

Use the `ctstopd` command to stop all Sun HPC ClusterTools software daemons on the cluster nodes. Once the Sun HPC ClusterTools 5 software is activated, `ctstopd` is available in `/opt/SUNWhpc/sbin`.

See TABLE 4-6 for a summary of the `ctstopd` options.

TABLE 4-6 `ctstopd` Options

Options	Description
General	
`-h`	Command help.
`-l`	Execute the command on the local node only.
`-R`	Specify the full path to be used as the root path.
`-x`	Turn on command debug at the specified nodes.
Centralized Operations Only	
`-g`	Generate node lists of successful and unsuccessful `ctstopd` operations.
`-k`	Specify a central location for storing copies of local log files.
`-n`	List of nodes where daemons will be stopped. Separate names by comma.
`-N`	File containing list of nodes where daemons will be stopped. One node per line.
`-r`	Remote connection method: rsh, ssh, or telnet.
`-S`	Specify full path to an alternate ssh executable.

Stopping Daemons From a Central Host

This section shows how to stop Sun HPC ClusterTools software daemons from a central host.

Stop Daemons on Specified Cluster Nodes

CODE EXAMPLE 4-28

```
# ./ctstopd –N /tmp/nodelist –r rsh
```

CODE EXAMPLE 4-28 stops the Sun HPC ClusterTools software daemons on the nodes listed in `/tmp/nodelist`. The remote connection method is rsh.

CODE EXAMPLE 4-29

```
# ./ctstopd –N /tmp/nodelist –r rsh –k /tmp/cluster-logs –g
```

CODE EXAMPLE 4-29 is the same as CODE EXAMPLE 4-28, except it specifies the options –k and –g to gather log information centrally and to generate pass and fail node lists.

Stopping Daemons on the Local Node

Stop Daemons Locally

CODE EXAMPLE 4-30

```
# ./ctstopd –l
```

CODE EXAMPLE 4-30 stops the Sun HPC ClusterTools software daemons on the local node.

CHAPTER 5

Set Up an NFS Server for Sun HPC ClusterTools Software Clients

This chapter explains how to use the `ctnfssvr` command to

- Install Sun HPC ClusterTools NFS server software.
- Enable Sun HPC ClusterTools NFS server software.
- Disable Sun HPC ClusterTools NFS server software.
- Remove Sun HPC ClusterTools NFS server software.

Overview

The following is a summary of the basic operations involved in installing Sun HPC ClusterTools software in an NFS client/server configuration:

- Use `ctnfssvr` to install the non-root Sun HPC ClusterTools software packages on the NFS server and to enable the Sun HPC ClusterTools software.
- Use `ctgui` or the CLI command `ctinstall -c` to install the root packages on the cluster nodes.
- If you will use the NFS server as a cluster node, use `ctgui` or `ctinstall -c` to install the root packages on the NFS server as well.

Note – You can perform these operations in any order, but the Sun HPC ClusterTools software can not be activated until the NFS server is installed, enabled, and properly mounted.

See Chapter 3 for instructions on using `ctgui` and Chapter 4 for instructions on using `ctinstall` and the other CLI commands. The rest of this chapter explains how to use `ctnfssvr` to install the non-root packages on the NFS server.

See TABLE 5-1 for a summary of the `ctnfssvr` options.

TABLE 5-1 `ctnfssvr` Options

Options	Description
Operations	
`-i`	Install non-root packages. The following options can be used only with `-i`.
	`-d` Specify a non-default *install from* location. The default is *distribution*`/hpc/Product`, relative to the directory where `ctnfssvr` is invoked.
	`-s` Specify a security option for Sun CRE. The choices are `sunhpc_rhosts`, `rhosts`, `des`, and `krb5`. The default is `sunhpc_rhosts`.
	`-t` Specify a non-default *install to* location. The default is `/export`.
`-r`	Remove the non-root packages.
`-D`	Disable serving Sun HPC ClusterTools software clients over NFS.
`-E`	Enable serving Sun HPC ClusterTools software clients over NFS.
Other	
`-h`	Command help.
`-x`	Turn on command debug mode.

Note – You must be logged in as superuser to run `ctnfssvr`.

Special Note for Activating Sun HPC ClusterTools Software in an NFS Configuration

When you want to activate Sun HPC ClusterTools software in an NFS configuration, you must ensure that the activation tool is able to locate the software mount point. You can do this in either of the following ways:

- Use the activation tool under the NFS mount point on an NFS client: *mount_point*`/*/SUNWhpc/HPC5.0/bin/Install_Utilities/bin`
- Or, explicitly mount the directory in which `SUNWhpc` is installed. For example, if `SUNWhpc` is installed in `/export/apps`, enter the following on an NFS client:

```
# mount server:/export/apps mount_point
```

Sun HPC ClusterTools software activation is described in Chapter 3 (for `ctgui`) and Chapter 4 (for `ctact`).

Installing the Non-Root Packages on an NFS Server

Install NFS Server Software

CODE EXAMPLE 5-1

```
# ./ctnfssvr –i
```

CODE EXAMPLE 5-1 sets up the NFS server and installs the non-root packages from the Sun HPC ClusterTools software suite. However, the NFS server has not been enabled to service Sun HPC ClusterTools software requests from the NFS clients (that is, the cluster nodes).

CODE EXAMPLE 5-2

```
# ./ctnfssvr –i –E
```

CODE EXAMPLE 5-2 is the same as CODE EXAMPLE 5-1, except that the `–E` option enables NFS services for the Sun HPC ClusterTools software clients.

Other `ctnfssvr` Examples

CODE EXAMPLE 5-3

```
# ./ctnfssvr –E
```

CODE EXAMPLE 5-3 enables the NFS server software after it has been installed.

CODE EXAMPLE 5-4

```
# ./ctnfssvr –D
```

CODE EXAMPLE 5-4 disables the NFS server software.

CODE EXAMPLE 5-5

```
# ./ctnfssvr -r
```

CODE EXAMPLE 5-5 removes all NFS server software. If the software is in the enabled state, `-r` will automatically disable it before the packages are removed.

CHAPTER 6

Additional Steps

This chapter describes the post-activation phase—that is, the final steps needed to get your Sun HPC system ready for use after the software has been activated.

Verifying Basic Functionality

Use the procedures in this section to test the cluster's ability to perform basic operations.

Note – You need to have `/opt/SUNWhpc/bin` in your path for many of the following procedures.

▼ To Display Information About Cluster Nodes

1. **Display information about the cluster nodes:**

   `% mpinfo –N`

   The following is an example of `mpinfo -N` output for a two-node system:

```
% mpinfo -N
NAME   UP  PARTITION  OS     OSREL  NCPU  FMEM   FSWP   LOAD1  LOAD5  LOAD15
node1  y   -          SunOS  5.8    1      7.17  74.76  0.03   0.04   0.05
node2  y   -          SunOS  5.8    1     34.70  38.09  0.06   0.02   0.02
```

2. **(Optional) Restart the node daemons.**

   If any nodes are missing from the list or do not have an entry in the `UP` column, restart their node daemons. See “Start Sun HPC ClusterTools Software Daemons” on page 42 for instructions on starting Sun CRE daemons.

   If Sun CRE daemons do not start, check `/var/adm/messages` for error messages.

   If an authentication error is logged in `/var/adm/messages`, the problem may be that the software was not activated on the cluster’s master node. The following is an example of this error message:

```
mpinfo: tmrte_rdb_do_find: tmrte_rdb_find_3() on node2: RPC call failed:
mrte_rdb_find_3: RPC: Authentication error:; why = Client credential too weak
```

   This will happen if you use the default authentication method, `sunhpc_rhosts`, and activate software from some node other than the master node. Activating the software from the master node causes the `sunhpc_rhosts` file to be automatically populated with a complete list of the nodes in the cluster. If the software is activated from a non-master node, the `sunhpc_rhosts` file is not created automatically. The `sunhpc_rhosts` method uses this list to determine which nodes are members of the cluster and are permitted to communicate with each other.

   Verify that the file `/etc/sunhpc_rhosts` exists on the master node and that it contains a list of all the nodes in the cluster. For example, in a cluster consisting of node1 and node2, the `sunhpc_rhosts` file might look like this:

```
node1
node2
```

▼ To Customize the Key Files

When executing programs, Sun CRE checks the credentials passed in each remote procedure call against the contents of a key file stored on each node.

The installation procedure creates key files that contain a default password. For security reasons, you should customize these files with your choice of cluster password immediately after installation. The password should consist of 10-20 alphanumeric characters.

- **As superuser, run the** `set_key` **script on each node of the cluster *and* on any nodes outside the cluster that may be accessed by a program running on the cluster.**

```
# /etc/opt/SUNWhpc/HPC5.0/etc/set_key
```

This script stores a password in `/etc/hpc_key.`*cluster_name*.

Alternatively, you can use the Cluster Console tools to update the key files on all the nodes from a single host.

Executing the `set_key` script on all the cluster nodes in this way assures that the same password is used across a cluster. See the *Sun HPC ClusterTools Software Administrator's Guide* for more information.

▼ To Create a Partition

The Sun CRE `mprun` command runs only within a Sun CRE *partition*, which is a logical set of nodes. This means you must create a partition containing the cluster nodes before you can execute an `mprun` operation.

Note – The following procedure is a one-time task. Perform it only once on a single cluster node for the entire cluster.

---

1. **Log in as superuser on any node in the cluster.**
2. **Run the** `part_initialize` **script on one of the cluster nodes.**

   Perform this step only one time on any cluster node.

```
# /opt/SUNWhpc/sbin/part_initialize
```

This script creates a partition named `all`, which includes all the nodes in the cluster. This partition can be used in subsequent verification tests.

For more information about partitions, see the *Sun HPC ClusterTools Software Administrator's Guide*.

▼ To Verify Sun CRE Setup

- **After you have created the partition** `all`**, run** `mpinfo -N`**.**

  The output of `mpinfo -N` should show the nodes are in the partition, `all`.

```
% mpinfo -N
NAME   UP  PARTITION  OS     OSREL  NCPU  FMEM  FSWP   LOAD1  LOAD5  LOAD15
node1  y   all        SunOS  5.8    1     8.26  74.68  0.00   0.01   0.03
node2  y   all        SunOS  5.8    1    34.69  38.08  0.00   0.00   0.01
```

▼ To Check Host Names

- **Invoke the** `mprun hostname` **utility.**

  This should display all the host names in your cluster, printing them one per line. The following example illustrates this output in a cluster that has two nodes:

```
% mprun –Ns –np 0 hostname
node1
node2
```

▼ To Verify That Sun CRE Executes Jobs

- **Run the following test:**

```
% mprun –np 0 uname –a
SunOS  node1 5.8 Generic sun4u sparc  SUNW,Ultra-5_10
SunOS  node2 5.8 Generic sun4u sparc  SUNW,Ultra-5_10
```

Enabling Close Integration With Batch Processing Systems

Sun CRE provides close integration with several distributed resource management systems. For information on how that integration works and how to set up the integration for each of the supported resource managers, refer to the *Sun HPC ClusterTools Software Administrator's Guide*.

Verifying MPI Functionality

This section explains how to verify that the appropriate network interfaces are available and how to test MPI communications.

Verifying Network Interface

The communication protocol to be used must be listed in the configuration file `hpc.conf`, and, for internode communications, associated with the appropriate network interface(s).

The default `hpc.conf` file provided with Sun HPC ClusterTools software includes the most commonly used configurations.

The `hpc.conf` file lists the three communication protocols supplied with the software: SHM (shared memory), RSM (remote shared memory), and TCP (Transport Control Protocol). The entry in the `LIBRARY` column, `()`, indicates that the protocol modules are installed in the default location.

```
# List the available Protocol Modules
# PMODULE LIBRARY
Begin PMODULES
shm       ()
rsm       ()
tcp       ()
End PMODULES
```

In addition, the `hpc.conf` file associates each protocol module with one or more types of network interface. The RSM protocol is associated, by default, with all interfaces to the Sun Fire high-performance interconnect (`wrsm`):

```
# RSM settings
# NAME  RANK    AVAIL
Begin PM=rsm
wrsm    20      1
End
```

The TCP protocol is associated with a large number of interface types. These are listed in the `hpc.conf` template:

```
idn – 16k   (StarFire Inter-Domain Network)
scid – 32K  (Dolphin SCI)
ba – 8K     (Sun ATM)
fa – 8K     (Fore ATM(SPANS))
acip – 8K   (Adaptec ATM)
anfc – 16K  (Ancor Fibre Channel)
bf – 4K     (Branch FDDI)
be – 4K     (SPARC Ethernet 100mbit)
hme – 4K    (SPARC Ethernet 100mbit)
le – 4K     (SPARC Ethernet 10mbit)
smc – 4K    (SMC Ethernet 10mbit)
```

Note – Inclusion of any network interface in this file does not imply that Sun Microsystems supports that network interface in a Sun environment.

If the network interface you use for TCP communication is not among those listed in `hpc.conf`, you must add it and then restart Sun CRE.

▼ To Add a TCP Interface Type

1. **Decide upon a rank value.**

   The rank indicates the relative preference of that interface compared with others that are available, with the lowest rank most preferred.

2. **Add the interface name and rank value to** `hpc.conf` **in the** `PM=tcp` **section:**

```
TCP Settings
NAME    RANK    MTU     STRIPE LATENCY  BANDWIDTH
Begin PM=tcp
midn    0       16384   0      20       150
idn     10      16384   0      20       150
...
End PM
```

The MTU, STRIPE, LATENCY, and BANDWIDTH columns are placeholders whose values are not used at this time. Simply repeat the values shown for the other TCP-enabled interfaces (16384, 0, 20, and 150).

For example, you could add the following entry to the `hpc.conf` file to include an interface named `niki` with a preference ranking of 50:

```
niki     50   16384    0      20       150
```

3. **Restart Sun CRE.**

Verifying MPI Communications

You can verify MPI communications by running a simple MPI program.

▼ To Verify MPI Communications

1. **Ensure that one of the supported compilers is installed on your system.**

See “Supported Compilers” on page 15.

2. **Run one of the sample MPI programs**

Three simple Sun MPI sample programs are available in the directory `/opt/SUNWhpc/examples/mpi`:

- `connectivity.c` – A C program that checks the connectivity among all processes and prints a message when it finishes.
- `monte.f` – A Fortran program that involves each MPI process in calculating an estimate of π using a Monte-Carlo method.
- `prime.cc` – A C++ program that calls each non-root rank to send a list of numbers to root. Root checks the incoming lists for prime numbers and generates a report of the prime numbers it finds.

See the `Readme` file in the same directory for instructions on how to use the examples. The directory also contains a make file, `Makefile`. The full text of both code examples is also included in the *Sun MPI Software Programming and Reference Guide*.

APPENDIX **A**

Required File and Directory Settings

Sun HPC ClusterTools software installation requires special settings on several files and file systems. TABLE A-1 describes these settings.

TABLE A-1 Installation Requirements

Files or Directories	Owner	Permissions	Comments
`/opt/SUNWhpc/hpc.conf`	superuser	644	Must be accessible by all nodes.
`/etc/sunhpc_rhosts`	—	600	—
CD-ROM *mount point*	—	—	Must be readable by superuser and accessible through a common path from all nodes.
NFS client mount point	—	—	Server mount point must exist on all client nodes.

APPENDIX **B**

Packages

Sun HPC ClusterTools software includes numerous packages. Their individual sizes are listed in TABLE B-1. For non-NFS configurations, the total space requirement per node is approximately 92 Mbytes. In an NFS configuration, the packages installed on each NFS client node occupy only 146 Kbytes.

TABLE B-1 Sun HPC ClusterTools Software Packages

Package Name	Package Size (Kbytes)
SUNWcre	2809
SUNWcremn	112
SUNWcrert (root package)	55
SUNWcrex	1900
SUNWctprf	8203
SUNWhpamn	64
SUNWhpcat	1048
SUNWhpctm	88
SUNWhpmmn	27
SUNWhpmsc	20
SUNWmpi	10841
SUNWmpimn	1099
SUNWmpirt	43
SUNWmpix	11297
SUNWpfs*	13
SUNWprfmn	29
SUNWprism	12185

TABLE B-1 Sun HPC ClusterTools Software Packages *(Continued)*

Package Name	Package Size (Kbytes)
SUNWprsmn	903
SUNWprsmx	14859
SUNWs3l	13946
SUNWs3lmn	794
SUNWs3lx	11668

* Although Sun HPC ClusterTools software includes a PFS package, that package provides warning messages only. Those warning messages alert users that PFS executables from Sun HPC ClusterTools 4 software are not supported by Sun HPC ClusterTools 5 software. *Sun PFS is no longer provided as part of Sun HPC Cluster Tools software.*