Sun HPC ClusterTools™ 3.1 Installation Guide

THE NETWORK IS THE COMPUTER™

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303-4900 USA
650 960-1300 Fax 650 969-9131

Part No. 806-3730-10
March 2000, Revision A

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

Please
Recycle

Contents

Preface

This document describes the procedure for installing Sun HPC ClusterTools™ 3.1 software.

These instructions are intended for an experienced system administrator. For example, to install the ClusterTools software on one or more nodes, the installer must be familiar with the following tasks in a Solaris 2.6, Solaris 7, or Solaris 8 environment:

- Logging in as root.
- 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 debug information.
- Starting and stopping license daemons.
- Exporting and mounting an NFS file system and using commands and scripts, such as: `mount`, `share`, `/etc/init.d/nfs.server`.
- Enabling root login access to a server.
- Setting directory and file permissions to allow read and write access.

If Your Cluster Includes an SCI Network

If your cluster includes an SCI network, refer to the *Sun HPC SCI 3.1 Guide* for instructions on installing the SCI software packages and configuring the network.

Before You Read This Book

In order 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 3.1 Product Notes*
- Documentation that accompanied your Sun Enterprise™ server
- Documentation for the applicable Solaris operating environment
- Documentation for Platform Computing Corporation's LSF suite (if you will be using LSF as the cluster resource manager)

Note – If you will be using LSF as the cluster resource manager, the LSF software must be installed *before* you install the ClusterTools software.

Using UNIX Commands

This document may 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 2.x Handbook for SMCC Peripherals*
- AnswerBook2™ online documentation for the Solaris™ software environment
- Other software documentation that you received with your system

Typographic Conventions

TABLE P-1 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	Read Chapter 6 in the *User's Guide*. These are called *class* options. You *must* be superuser to do this.
AaBbCc123	Book titles, new words or terms, words to be emphasized	Read Chapter 6 in the *User's Guide*. These are called *class* options. You *must* be superuser to do this.
	Command-line variable; replace with a real name or value	To delete a file, type `rm` *filename*.

Shell Prompts

TABLE P-2 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

TABLE P-3 Related Documentation

Application	Title	Part Number
Sun HPC ClusterTools software	*Sun HPC ClusterTools 3.1 Product Notes*	806-4182-10
Administering Sun HPC ClusterTools software	*Sun HPC ClusterTools 3.1 Administrator's Guide*	806-3731-10
Installing and configuring SCI network	*Sun HPC SCI 3.1 Guide*	806-4183-10

For information about the LSF suite of resource management software, see the LSF documentation published by Platform Computing Corporation.

CHAPTER 1

Overview

FIGURE 1-1 shows an overview of the installation process using the Sun HPC ClusterTools software configuration tool, `install_gui`.

The configuration tool allows you to perform the following tasks:

- Install the ClusterTools 3.1 software
- Remove a specified version of installed ClusterTools software, either version 3.0 or 3.1
- Specify one version of ClusterTools software to be active, either version 3.0 or 3.1

Appendix A lists error messages and associated diagnostic tips. Appendix B contains a summary list of required file and directory settings.

Note – Appendix B of the *Sun HPC ClusterTools 3.1 Administrator's Guide* explains how to install the ClusterTools software using a command-line script in place of `install_gui`.

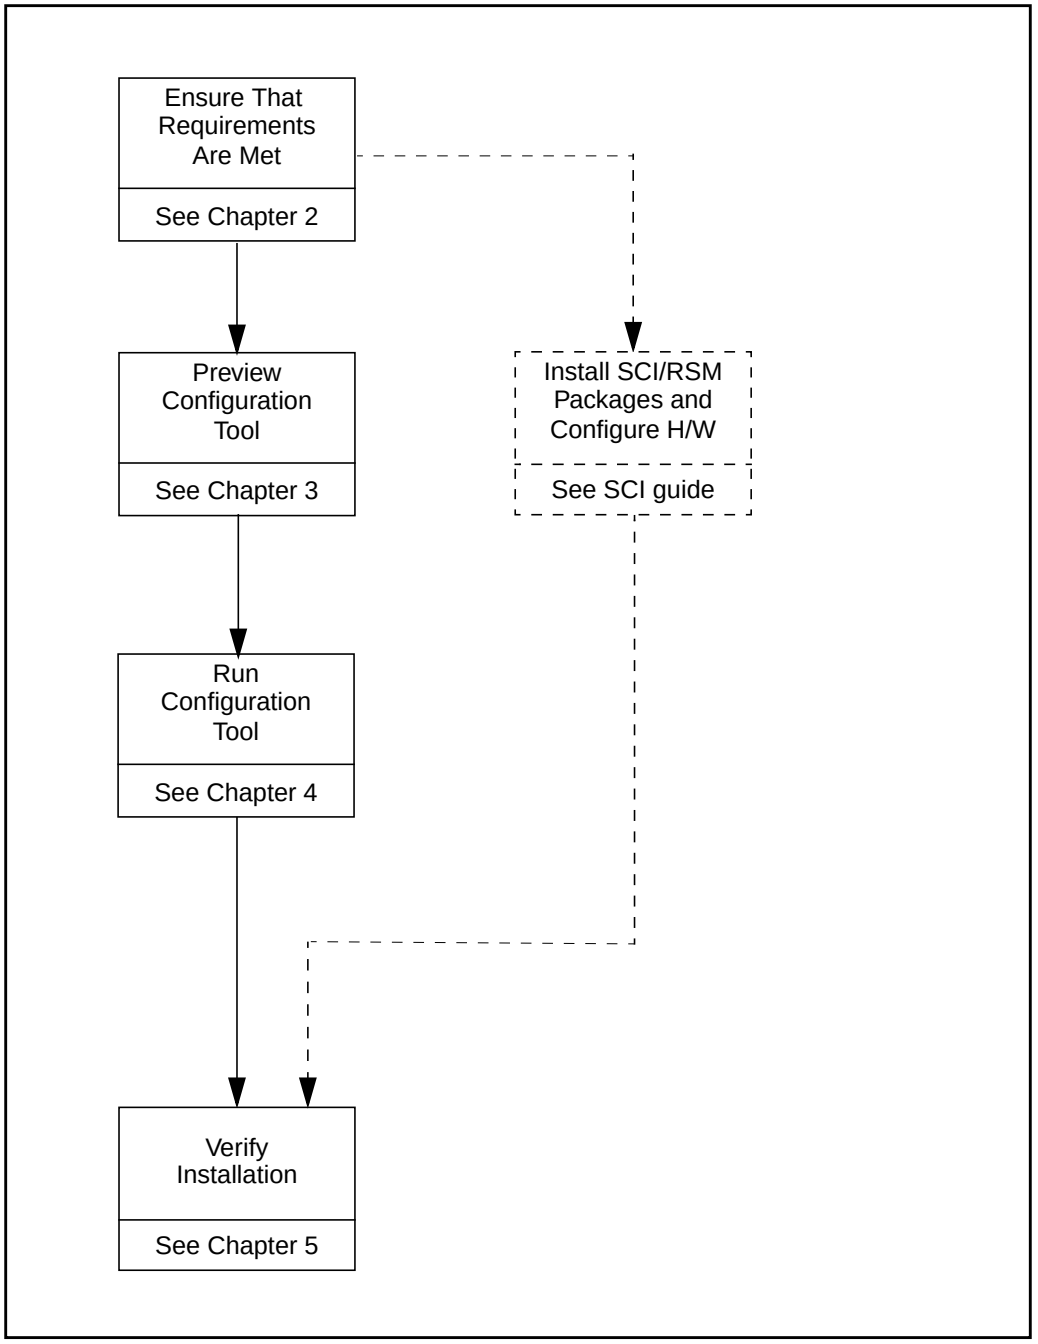

FIGURE 1-1 Overview of Sun HPC ClusterTools Software Installation

CHAPTER 2

Requirements for Installation

Before installing Sun HPC ClusterTools™ 3.1 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:

- Standard Requirements – For every cluster running Sun HPC ClusterTools 3.1 software.
- Custom Requirements and Recommendations – For selecting your choice of resource management software and for running the Sun Parallel File System (PFS) or Sun Scalable Scientific Subroutine Library (Sun S3L).
- Supported Compilers – For programming with Sun HPC ClusterTools 3.1 software.

Note – During installation, several files and directories will require specific permissions, ownership, or other settings. For a quick summary of information about several important settings, see Appendix B.

Standard Requirements

Before installing Sun HPC ClusterTools software, be sure that all hardware is installed and configured on a network (such as Ethernet). Any other connections, such as ATM or a terminal concentrator, must also be fully installed and configured.

TABLE 2-1 outlines additional prerequisites for installation.

TABLE 2-1 Standard Installation Prerequisites for Sun HPC ClusterTools 3.1 Software

Description	Requirement
Hardware	Sun UltraSPARC-based systems.
Disk space[1]	Approximately 85 Mbytes per node (This estimate does not include space required for AnswerBook2.)
Operating[2] environment	Solaris 2.6, Solaris 7, or Solaris 8
Resource manager	Your cluster must have as its resource manager either CRE 1.1, LSF 3.2.3 or LSF 3.2.4. The CRE packages will be installed with the rest of the ClusterTools 3.1 software. However, if you will be using LSF, it must be installed on the cluster before you install the ClusterTools software.
`/etc/system` file settings	Edit the `/etc/system` file on every node in your cluster. Add the following entries to the `/etc/system` file on each node: set pt_cnt=1024 set npty=1024 set sadcnt=2048 set nautopush=1024 set dosynctodr=0 set shmsys:shminfo_shmmax=2120000000 set shmsys:shminfo_shmmni=200 set shmsys:shminfo_shmseg=200 See the applicable Solaris administration documentation for more information about `/etc/system`.
Earlier versions of Sun HPC Software	Sun HPC 2.0 Software must be removed. Sun HPC ClusterTools 3.0 software does *not* need to be removed. ClusterTools versions 3.0 and 3.1 can coexist on a cluster.

1. If the ClusterTools software is installed on an NFS server, this disk storage allowance will be needed only on the NFS server. If the software is installed locally, each node must have 85-Mbytes available for the ClusterTools installation.
2. If you install Sun HPC ClusterTools 3.1 software locally on nodes that have the Solaris 2.6 operating environment, only 32-bit ClusterTools will be installed. If you later upgrade the nodes to Solaris 7 or Solaris 8, you will have to install the ClusterTools 3.1 software again for the 64-bit libraries to be installed. A second installation is not required when the Sun HPC ClusterTools 3.1 software is installed on an NFS server. NFS-mode installations always install support for Both 32-bit and 64-bit libraries.

Custom Requirements and Recommendations

Choosing Resource Management Software

The Sun HPC ClusterTools 3.1 software distribution includes the CRE 1.1 resource manager. No other resource management software is needed for interactively launching Sun MPI jobs. For launching jobs in batch mode, you can use the LSF 3.2.3 or 3.2.4 suite in place of CRE. If you do plan to use LSF, the requirements outlined in TABLE 2-2 apply.

TABLE 2-2 Requirements for Using LSF

Description	Requirement
Installing LSF	The LSF software must be installed before the ClusterTools software is installed.
Required LSF components	The following LSF components are required: • LSF Base • LSF Batch • LSF Parallel
Required LSF patch	Contact Platform Computing Corporation for the LSF patch that is required for compatibility with ClusterTools 3.1.
LSF cluster	Every node that will access Sun HPC ClusterTools software must also be in a corresponding LSF cluster. See the discussion of the `lsf.cluster.`*clustername* configuration file in the *LSF Batch Administrator's Guide* for additional information.

PFS and Sun S3L

TABLE 2-3 outlines prerequisites for installing the Sun PFS and Sun S3L components of the Sun HPC ClusterTools software.

TABLE 2-3 Prerequisites for Installing PFS and Sun S3L Software Components

Component	Requirement
PFS	If you intend to use only part of a disk partition for PFS file systems, you should take time to plan your file system before installing the software. See the *Sun HPC ClusterTools Administrator's Guide* for more information about PFS.
Sun S3L	Sun S3L requires Sun Performance Library™ to be installed. Sun Performance Library is included with the WorkShop Compilers Fortran and Performance WorkShop Fortran compilers, v4.2 and 5.0.

Supported Compilers

The following compilers are supported with Sun HPC ClusterTools 3.1 software:

- Sun WorkShop Compilers C/C++ 4.2
  (also included in Sun Visual WorkShop C++ 3.0)
- Sun WorkShop Compilers Fortran 4.2
  (also included in Sun Performance WorkShop Fortran 3.0)
- Sun Visual WorkShop C++ 5.0
- Sun Performance WorkShop Fortran 5.0

The compilers can be installed either before or after the ClusterTools software.

CHAPTER 3

Preview of the Configuration and Installation Tool, `install_gui`

General

Sun HPC ClusterTools 3.1 software includes the configuration and installation tool `install_gui`, which simplifies the ClusterTools software installation process. This tool provides a graphical user interface (GUI) through which you enter information about the cluster configuration and how you want the ClusterTools software to be installed.

The configuration tool saves the information you enter in a configuration file and then initiates the installation process. The installation is guided by the contents of the configuration file. You are asked to specify the name and location of the configuration file; the default file name will be `hpc_config`.

The configuration file will also be referenced any time you make subsequent changes to the ClusterTools installation status, such as adding or removing nodes from the cluster.

Note – There are also four `install_gui` command options that may be of interest to some users. They are described in the section “`install_gui` Usage” on page 8.

ClusterTools 3.1 can coexist with ClusterTools 3.0. If ClusterTools 3.0 is already installed on the target node(s), you can install ClusterTools 3.1 without removing 3.0. However, when a cluster contains both versions of ClusterTools, you must explicitly deactivate one version and activate the other using controls provided by the `install_gui` tool.

This chapter discusses the various kinds of information the configuration tool will ask you to supply. Knowing about these topics in advance is likely to speed up the installation process. It begins with a discussion of the `install_gui` command usage.

Note – If, for some reason, you can't use `install_gui` to perform the installation, you can set up the `hpc_config` file using a text editor and run the installation scripts from a shell prompt. See Appendix B of the *Sun HPC ClusterTools 3.1 Administrator's Guide* for instructions on how to install ClusterTools software without running `install_gui`.

`install_gui` Usage

The `install_gui` is supplied on the CD-ROM that contains the ClusterTools 3.1 software. Its path is `/cdrom/hpc_3_1_ct/Product/Install_Utilities`.

The `install_gui` command can be used by itself—that is, without switches or command-line arguments, as follows:

```
# install_gui
```

There are, however, four command-line options that can be used to achieve particular results. These are described in the following sections.

Specifying a Configuration File

You can specify a configuration file to be loaded automatically when the GUI is launched. If you want to load the default file, `hpc_config`, in this way, supply the name of the directory that contains `hpc_config` as an argument to the `-c` option. For example, if the `hpc_config` file is located in `/usr/admin`, enter

```
# install_gui -c /usr/admin
```

If the configuration file has some file name other than `hpc_config`, you must supply the file's full path name. You might choose to assign a nondefault name to a configuration file so that you can have more than one such file, each tailored to configure a ClusterTools software installation in a particular way.

For example, you might have a file customized for NFS installations on a four-node cluster and have named it `nfs-4`. If this configuration file is located in `/usr/admin`, you would load it from the command line by entering:

```
# install_gui -c /usr/admin/nfs-4
```

Alternatively, you can load a configuration file after the installation GUI is launched. To do this, select the Load option from the File pull-down menu and then specify the configuration file to be loaded. Again, specify either the directory containing `hpc_config` or the full path name of a configuration file that has a custom file name.

Specifying Custom rsh and telnet Executables

You can use the `-rshPATH` and `-telnetPATH` options to access nondefault rsh and telnet executables. For example, if the ClusterTools configuration will be using Kerberos software, you might want to specify an rsh executable that is customized for Kerberos-based authentication. For example, supply the full path to the desired rsh executable as an argument to the `-rshPATH:` option.

```
# install_gui –rshPath:/usr/krb5/bin/rsh
```

Specifying an Installation Stripe Size

The `-parallel:`*number* option can be used to specify the maximum size of the installation stripe—that is, the largest number of nodes that will be installed in parallel. The *number* argument must be <= 8. The following example limits the parallel installation stripe to 4.

```
install_gui –parallel:4
```

The parallel installation stripe defaults to 8.

Specifying the Resource Manager

Early in the installation process, the configuration tool will ask which resource manager will be used, CRE or LSF. If LSF, the LSF suite must already be installed before the ClusterTools software can be installed.

Choosing the Type of Installation

The configuration tool will also ask you to specify the type of installation to be performed, presenting the following as your options:

- Install locally on a single node – A single copy of the ClusterTools software will be installed on a single node.
- Install locally on a cluster of multiple nodes – A copy of the ClusterTools software will be installed locally on each node in the cluster.
- Install on an NFS server and mount remotely on other nodes in the cluster – A single copy of the ClusterTools software will be installed on an NFS server and remotely mounted on all the nodes in the cluster.

Directories Needed for Installation

The installation process will require the following directories:

- CD-ROM mount point – The CD-ROM must be mounted on all nodes on which the ClusterTools software will be installed.
- Configuration file directory – The installation process creates a configuration file with the default name `hpc_config`. You will be asked to specify where you want this file to be saved. You can also save it with a file name other than `hpc_config`. All nodes involved in the ClusterTools installation must have read and write access to the saved configuration file and to the directory in which it resides.
- Installation target directory – You will be asked to specify where you want the ClusterTools packages to be installed. If the software will be installed on an NFS server, the NFS installation directory must be mounted on each NFS client in the cluster. For local installations, the default installation directory is `/opt`. For NFS installations, specify an appropriate installation directory and NFS mount point.

Choosing an Installation Method

If you will be installing on more than one node (either locally or remotely from an NFS server), you will need to select between two methods for propagating the installation to all the nodes in the cluster. Your choices are

- *telnet* – If you choose the `telnet` option, you will have to supply the root password for every host in the cluster.
- *rsh* – If you choose `rsh`, all the nodes in the cluster must be made trusted hosts, at least during installation. You must also have permission to use `rsh` as root to all of the nodes in the cluster and, in case of an NFS installation, to the file server.

Supplying Host Names of Cluster Nodes

If installing locally on more than one node or via an NFS server, the configuration tool will ask you to supply the host names of the nodes on which the ClusterTools software will be installed. You can supply these names in either of two ways:

- Use the editor supplied by `install_gui` – The configuration tool provides an editing window for adding host names to and removing them from a list that it displays.
- Load a host file – The configuration tool allows you to load a file containing a list of host names. Each host name must be on a separate line. See FIGURE 3-1 for examples of the four possible host name formats. The format differences depend on two conditions:.
  - Whether you specify `telnet` or `rsh` for the installation method. If `telnet`, the node's password must be supplied (preceded by a colon).
  - Whether the nodes are connected to a terminal concentrator. If so, the host name must be followed by the host name of the terminal concentrator and the port ID to which that node is connected. The node host name, terminal concentrator host name and port ID are separated by forward slashes. See FIGURE 3-1 for examples of the four possible host name formats.

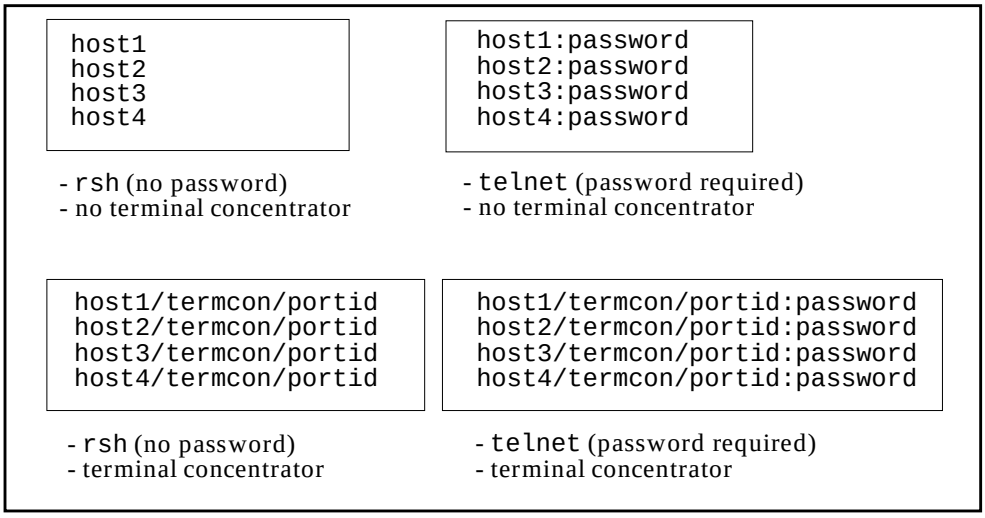

FIGURE 3-1 Examples of Host Name Files for Loading During Installation

---

Note – If you are installing the ClusterTools software locally on a single node, you must be logged in to the node on which the software will be installed. Consequently, you do not need to supply the node's host name.

Enable Root Login

You must be logged in as root to run the configuration tool. Since the default for most systems is to allow logins by root only on their console devices, you may need to edit the `/etc/default/login` file on each node to permit login by root. On each node, find the following line in the `login` file

```
CONSOLE=/dev/console
```

and add a # at the beginning of the line to convert it to a comment.

```
#CONSOLE=/dev/console
```

If your site's security guidelines require it, you should disable root login access after it's no longer needed by removing the comment character.

Choosing an Authentication Method

Authentication software provides increased levels of security, guarding against access by unauthorized users or programs. CRE supports two forms of authentication: Data Encryption Standard (DES) and Kerberos version 5. The configuration tool will ask you to select one of these methods or specify that you want no authentication method installed.

If you specify no authentication method, you will have the option of installing the file `sunhpc_rhosts`, which contains a list of the hosts that can access the cluster. If you to decline to install `sunhpc_rhosts`, your `.rhosts` file will be used for authentication instead.

Choosing a Master Node

The CRE consists of a set of daemons. A subset of these daemons, called the master daemons, run on a single node which you specify to be the master node. If you are installing the ClusterTools software on a single node, it will automatically be the master node.

Specifying the LSF Cluster Name (LSF Only)

If the LSF suite will be your resource manager, the configuration tool will ask for the LSF cluster name. This is the name assigned to the `LSF_CLUSTER_NAME` parameter in the `lsf_config` file. See the LSF installation documentation for information about this file.

LSF Parameter Modification (LSF Only)

If the LSF suite will be your resource manager, the configuration tool will ask if you want to modify certain LSF parameters that optimize HPC job launches. For more information about the LSF parameters that would be changed if you were to say yes, see “Configuring For Fast Interactive Batch Response Time” in Appendix C of the *Sun HPC ClusterTools 3.1 Administrator’s Guide*.

CHAPTER 4

Using install_gui

This chapter shows how to use the configuration and installation tool `install_gui` to perform the following tasks:

- Install the ClusterTools 3.1 software packages.
- Remove either the ClusterTools 3.0 or 3.1 software packages.
- Select one ClusterTools version, either 3.0 or 3.1, to be the active version.

These descriptions are preceded by instructions on how to start up `install_gui`.

Note – You must be logged in as root to run `install_gui`.

Starting install_gui

To start the configuration/installation tool

1. **Mount the CD-ROM path on all the nodes in the cluster.**
2. **Load the CD-ROM containing the ClusterTools 3.1 software in the CD-ROM drawer.**
3. **Log in to one of the nodes in the cluster as root and type the path for the `install_gui` command.**

```
# /cdrom/hpc_3_1_ct/Product/Install_Utilities/install_gui
```

When `install_gui` starts up, it displays a panel that asks you to choose one of three tasks:

- Install ClusterTools 3.1 software

- Remove installed ClusterTools
- Set Active Version

These tasks are respectively summarized in the next three sections.

If you loaded a configuration file when you launched `install_gui`, all three tasks will be active—that is, can be selected. However, if no configuration file has been loaded yet, you will be able to select only the Install ClusterTools 3.1 task. To load a configuration file from within the GUI, pull down the File menu and select Load. Then specify the applicable configuration file via the resulting dialog window.

Select the desired task and then click on the Next button to continue.

Note – For detailed explanations of how to use `install_gui`, pull down the Help menu on any GUI panel and select the Help with Configuration Tool item.

Installing ClusterTools 3.1

When you select the Install ClusterTools 3.1 task, `install_gui` will display a sequence of panels. The first three panels ask for various kinds of information about the cluster configuration and the kind of installation to be performed. Supply the information requested by each panel, clicking on the Next button to advance to the next panel.

The last panel provides a summary of the information you supplied on the previous panels as well as Install and Exit buttons. If the summary is accurate, you can do either of the following:

- Click on Exit to leave the `install_gui` environment. If you choose this option, you will be given the opportunity to save the information you entered in a configuration file of your choice. If you choose to exit and save your configuration, you can use that file to define the cluster configuration for subsequent operations on the ClusterTools software.
- Click on the Install button to start the installation process.

Make your choice and click on Save. This will start the installation process. An installation status window will then appear and report on the installation progress.

When the installation is complete, you will be given the choice of displaying the installation log file or saving it.

Note – If an error is detected during installation, the event will be logged. Some errors will only be reported and installation of the software will continue. More serious errors will stop the installation. See Appendix A “Installation Error Messages” for a list of error codes and recommended corrective action.

Removing ClusterTools 3.0 or 3.1

When you select the Remove Installed ClusterTools task, `install_gui` will display a panel that requests the following information:

- Choice of ClusterTools version to be removed, 3.0 or 3.1.
- The full path of the software removal script. The default path is `/cdrom/hpc_3_1_ct/`.

When you have supplied this information, click on the Remove button to start the software removal process. You will be asked to specify which file it should use as the configuration definition for the removal.

Make your choice and click on Save. This will start the removal process. A removal status window will then appear and report on the removal progress.

Note – If an error is detected during software removal, the event will be logged. Some errors will only be reported and removal of the software will continue. More serious errors will stop the removal process.

Selecting an Active ClusterTools Version

When both ClusterTools versions (3.0 and 3.1) are installed, you must specify which version is active and which is inactive. You do this using the `install_gui` select-version panel. There are two ways that you will encounter this panel:

- If ClusterTools 3.0 is present on the nodes when ClusterTools 3.1 is installed, the select-version panel will automatically appear right after installation of the 3.1 software completes.
- You can bring up the select-version panel by highlighting Set Active Version on the `install_gui` initial panel.

For example, if both ClusterTools versions are installed and you want to remove the active version, do the following:

1. Go to the initial panel of the `install_gui`, select the Set Active Version task, and click on the Next button. The select-version panel will appear.
2. Specify the full path of the script that will be used to deactivate/activate the versions. The default path is `/cdrom/hpc_3_1_ct/`.
3. In the deactivate section of the select-version panel, specify the version to be deactivated and click on the Deactivate button.
4. In the activate section of the panel, specify the version to be activated and click on the Activate button.

A warning will appear indicating that daemons will be shut down; you will be asked if you wish to continue. If you click on Yes the version selection script will start.

Note – When both versions of ClusterTools are resident on a cluster, you cannot remove the active version. You must first deactivate it and then activate the version that is to remain. If only one version of ClusterTools is on a cluster, you can remove it without first deactivating it.

The select-version controls can also be used to switch back and forth between ClusterTools versions even when you aren't planning to remove a version.

CHAPTER 5

Finishing Up

This chapter describes the post-installation phase—the final steps needed to get your Sun HPC system ready for use. These steps are organized into two main categories:

- The steps described in the first section apply only to clusters using the CRE resource manager.
- The steps described in the second section apply equally to clusters using either resource management software, LSF or CRE.

If your cluster uses the LSF suite for resource management, the LSF software must have been installed and verified before the ClusterTools software was installed. Refer to the *LSF Batch Administrator's Guide* for guidance in setting up and verifying LSF clusters.

Note – If your cluster will include an SCI network, its installation and configuration must be complete before proceeding with the finishing up phase. See the *Sun HPC SCI 3.1 Guide* for network configuration and software installation instructions. That document also provides tips on troubleshooting the SCI network.

Verify Basic Functionality (CRE Only)

Use the procedures described 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.

Run `mpinfo`

Run `mpinfo -N` to display information about the cluster nodes. 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
host1  y   -          SunOS  5.6    1      7.17  74.76  0.03   0.04   0.05
host2  y   -          SunOS  5.6    1     34.70  38.09  0.06   0.02   0.02
```

If any nodes are missing from the list or do not have a `y` entry in the `UP` column, restart their node daemons. See the *Sun HPC ClusterTools Administrator's Guide* for instructions on starting node daemons.

Create the `all` Partition

The CRE's `mprun` command will run only within a CRE *partition*. Log in as root on any node in the cluster and run the `part_initialize` script. This will create a partition named `all`, consisting of all the nodes in the cluster, which can be used in subsequent verification tests.

```
# /opt/SUNWhpc/etc/part_initialize
```

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

Verifying the CRE Setup

After you have configured the CRE into the partition `all`, run `mpinfo -N` again. This time, its output should show the nodes are in the partition `all`.

```
% mpinfo -N
NAME   UP  PARTITION  OS     OSREL  NCPU  FMEM   FSWP   LOAD1  LOAD5  LOAD15
host1  y   all        SunOS  5.6    1      8.26  74.68  0.00   0.01   0.03
host2  y   all        SunOS  5.6    1     34.69  38.08  0.00   0.00   0.01
```

Run `mprun -Ns`

Test `mprun` by launching the `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
host1
host2
```

Verify That the CRE Executes Tasks

To verify that the CRE is able to execute a simple program, run the following test:

```
% mprun –np 0 uname –a
```

Verify MPI Functionality (CRE and LSF)

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

If you will be using the Sun Parallel File System (PFS) software, refer to the *Sun HPC ClusterTools Administrator's Guide* for more information about configuring PFS file systems.

Verify Network Interface

If the network interface you are using is not listed in the `hpc.conf` file, you will need to add it to the `Netif` section of the `hpc.conf` file and restart your resource manager (LSF or CRE). If the interface is not added to this file, MPI communications may not work.

```
Begin Netif
NAME    RANK    MTU     STRIPE  PROTOCOL LATENCY  BANDWIDTH
midn    0       16384   0       tcp      20       150
idn     10      16384   0       tcp      20       150
sci     20      32768   1       tcp      20       150
mscid   30      32768   1       tcp      20       150
scid    40      32768   1       tcp      20       150
scirsm  45      32768   1       rsm      20       150
mba     50      8192    0       tcp      20       150
ba      60      8192    0       tcp      20       150
mfa     70      8192    0       tcp      20       150
fa      80      8192    0       tcp      20       150
macip   90      8192    0       tcp      20       150
acip    100     8192    0       tcp      20       150
manfc   110     16384   0       tcp      20       150
anfc    120     16384   0       tcp      20       150
mbf     130     4096    0       tcp      20       150
bf      140     4096    0       tcp      20       150
mbe     150     4096    0       tcp      20       150
be      160     4096    0       tcp      20       150
mqfe    163     4096    0       tcp      20       150
qfe     167     4096    0       tcp      20       150
mhme    170     4096    0       tcp      20       150
hme     180     4096    0       tcp      20       150
mle     190     4096    0       tcp      20       150
le      200     4096    0       tcp      20       150
msmc    210     4096    0       tcp      20       150
smc     220     4096    0       tcp      20       150
End Netif
```

This information can be acquired by using the Solaris `ifconfig` command. All that is required is the interface name, an `MTU` value, and a ranking. Interfaces with lower ranks are considered *preferred*; that is, they are chosen before interfaces with higher rank values.

For example, the following is an entry that could be made to the file to include a myrinet interface:

```
miri    162    16384   0      tcp      20      150
```

Verify MPI Communications

You can verify MPI communications by running a simple MPI program. To do so, you must have one of the supported compilers installed on your system. See “Supported Compilers” on page 6 of this manual for more information.

For information about running programs with LSF, see Platform Computing Corporation’s LSF documentation. For information about running programs with the CRE, see the *Sun HPC ClusterTools 3.1 User’s Guide*.

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

- `connectivity.c` – This is a C program that checks the connectivity among all processes and prints a message when it finishes.
- `monte.f` – This is a Fortran program that involves each MPI process in calculating an estimate of π using a Monte-Carlo method

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 Chapter 3 of the *Sun MPI Programming and Reference Guide*.

APPENDIX **A**

Installation Error Messages

During the installation process, output from `hpc_install` is written to a log file:

`/var/log/HPC-install.`*hostname*

TABLE A-1 shows each installation error message, a description, examples, and hints on how to resolve the error. (See the section that follows the table for related notes.)

Many aspects of the Sun HPC ClusterTools software installation process are controlled by a configuration file called `hpc_config`. Some of these error messages refer to this file,which can be modified either with the installation tool (`install_gui`) or using a text editor. When you click on Install in the fourth panel of the installation tool, the tool runs a script, `hpc_install`. For more information about the `hpc_config` file and the `hpc_install` script, see the *Sun HPC ClusterTools Administrator's Guide*.

TABLE A-1 Error Messages

ErrorCode	Description
`ErrorCode 01:`	Unable to find the install scripts.
Hint:	Check the CD-ROM mount point.
`ErrorCode 02:`	A variable or entry has not been set in the `hpc_config` file.
Example:	`INSTALL_LOC is not set in the hpc_config file.`
Hint:	Open the `hpc_config` file with either the configuration tool or a text editor to examine it for an unset entry.
`ErrorCode 03:`	A variable or entry in the `hpc_config` file has an invalid or illegal value.
Example:	`Illegal value for MODIFY_LSF_PARAM (`*value*`).` `Illegal value for INSTALL_METHOD (`*install_method*`).`
Hint:	Examine the entry in the `hpc_config` file.
`ErrorCode 04:`	Wrong version of LSF is installed.

TABLE A-1 Error Messages *(Continued)*

ErrorCode	Description
Example:	`LSF version 3.2.3 or higher is needed for the HPC software to run correctly.`
Hint:	Run `/opt/SUNWlsf/bin/lsid -V` to verify that you have the correct version of LSF installed. If you have the correct version installed, the command should return something like this: `LSF 3.2.3,` *date*
ErrorCode 05:	Unable to access the given directory, no such directory.
Example:	`/opt: no such directory.` `/opt/SUNWhpc: not a directory.`
Hint:	Verify that the given directory exists.
ErrorCode 06:	Unable to find one of the install scripts.
Example:	`$0: Cannot find rdcfg script.`
Hint:	Check the CD-ROM mount point.
ErrorCode 07:	Unable to remove SYNC files.
ErrorCode 08:	Unable to write into `/opt`.
Example:	`/opt: not writable.`
Hint:	Check the permissions of the `/opt` directory.
ErrorCode 09:	Cannot have a space in a host name or cannot have multiple host names in this variable.
Example:	`Invalid NFS server name --` *servername*`.`
Hint:	Examine the `hpc_config` file. Verify that there is only one host name in the entry, and that the host name does not include a space.
ErrorCode 10:	You must be root to run the install scripts.
Example:	`You must be "root" to run hpc_install.`
Hint:	Execute `/usr/ucb/whoami`. If `whoami` does not return `root`, become root.
ErrorCode 11:	Invalid or unknown node.
Example:	*nodename*: `unknown node.`
Hint:	Verify that the node name is valid. `ping` the node name.
ErrorCode 12:	Unable to access the node.
Example:	*nodename*: `unreachable.`
Hint:	Verify that *nodename* is up. Attempt logging in to *nodename*.

TABLE A-1 Error Messages *(Continued)*

ErrorCode	Description
ErrorCode 13:	`/opt/SUNWhpc` is already linked to another directory. A previous version of the Sun HPC ClusterTools 3.1 software may already be installed. Unable to install current version.
Example:	*target* `is linked to /opt/SUNWhpc.`
Hint:	Verify that there is no other version of Sun HPC Software installed. If one does not exist, delete `/opt/SUNWhpc`.
ErrorCode 14:	The `/opt/SUNWhpc` directory already exists, possibly from a previously installed Sun HPC Software. Unable to install the current version of Sun HPC ClusterTools 3.1 software.
Example:	`/opt/SUNWhpc: directory already exists.`
ErrorCode 15:	Unable to find the install scripts.
Example:	`Error parsing install script directory name.`
Hint:	Check the CD-ROM mount point.
ErrorCode 16:	An older version of Sun HPC Software was installed, please remove it before installing the current version.
Example:	`Version` *version* `is installed; it does not match the version` *version* `from the CD-ROM: Please remove it.`
Hint:	To remove Sun HPC Software 2.0 or earlier, see the documentation that came with the earlier version of the software. To remove Sun HPC ClusterTools 3.0 software, see Chapter 4.
ErrorCode 17:	Insufficient amount of free disk space.
Example:	*Total–space–required* `free blocks of disk space necessary for installation;` *space_avail* `blocks available on` *directory–name.*
Hint:	Perform `df -k` on the directory specified in `INSTALL_LOC`. The owner should be the same as the user specified in `LSF_ADMIN`.
ErrorCode 18:	The path supplied as an installation location does not exist.
Example:	*Install–location* `does not exist.`
Hint:	Verify that the directory specified in *Install–location* exists.
ErrorCode 19:	A problem has occurred while trying to create a symbolic link from `/opt/SUNWhpc` to the path set in *pathname.*
Example:	`Unable to create link from /opt/SUNWhpc to` *pathname.*
Hint:	Verify that the *pathname* exists. Verify that `/opt` is writeable.
ErrorCode 20:	Unable to create files in `/tmp`.
Example:	`Unable to write /opt/HPC_SKIP_POSTINSTALL.`

TABLE A-1 Error Messages *(Continued)*

ErrorCode	Description
Hint:	Verify the permissions on `/tmp`. The directory must be writeable.
ErrorCode 21:	The package was not installed correctly. Install or reinstall it manually (as root).
Example:	`Unable to install` *package*. `Unable to reinstall` *package*.
Hint:	See the *Sun HPC ClusterTools 3.1 Administrator's Guide*.
ErrorCode 22:	Unable to execute the postinstall script.
Example:	`Error: Unable to execute` *filename*.
Hint:	Examine the CD-ROM mount point. Verify that the mount point is readable / accessible.
ErrorCode 23:	The package was not removed correctly. Remove it manually (as root).
Example:	`Unable to remove` *packagename*.
Hint:	See the *Sun HPC ClusterTools 3.1 Administrator's Guide*.
ErrorCode 24:	You have chosen the `smp-local` configuration, but you have listed more than one host name in the list of nodes.
Example:	`pathname/hpc_config: warning, more than one node for smp-local.`
Hint:	Either change the configuration type to `cluster-local` or remove all but one host name from the list of nodes.
ErrorCode 25:	The `preremove` script doesn't exist.
Example:	`Cannot find preremove script`, *filename*.
Hint:	Verify that `/opt/SUNWhpc/HPC3.1/etc` is accessible on that system.
ErrorCode 26:	Root is unable to write into the directory containing the `hpc_config` file.
Example:	`Unable to write into $CONFDIR. Please make it writeable to root.`
Hint:	Verify that the directory containing the `hpc_config` file is writeable by root.
ErrorCode 27:	The LSF software does not exist on the system.
Example:	`LSF is not installed, please install LSF version $LSF_RELEASE or greater before installing HPC3.1.`
Hint:	Verify that the LSF software is installed on the system.
ErrorCode 28:	The path set in the environment variable `LSF_ENVDIR` is invalid.

TABLE A-1 Error Messages *(Continued)*

ErrorCode	Description
Example:	`$CONFDIR/$CONF_FILE_NAME: illegal value for LSF_CONF_DIR ($LSF_CONF_DIR).`
Hint:	Verify in `/etc/lsf.conf`, that the path set for `LSF_ENVDIR` is valid
ErrorCode 29:	HPC 2.0 software exists on the system.
Example:	`/opt/SUNWhpc/HPC2.0 exists. Please remove HPC2.0 before installing HPC3.1.`
Hint:	Use the `hpc2.0` removal scripts to remove HPC 2.0 software.
ErrorCode 30:	Cannot set the node names in the `NODES` section of the `hpc_config` file.
Example:	`Unable to set the NODES entry. Please set the NODES variable in the hpc_config file.`
Hint:	Edit the `hpc_config` file and enter the node names in the `NODES` section.
ErrorCode 31:	Attempting to use the HPC 3.1 removal scripts to remove HPC 2.0.
Example:	`Unable to use the HPC3.1 hpc_remove script to remove HPC2.0. Use the hpc 2.0 removal script to remove HPC2.0 from this system.`
Hint:	Use the applicable removal script.
ErrorCode 32:	Root is unable to write into the directory containing the HPC config file.
Example:	`Unable to write into $CONFDIR. Please make it writeable to root.`
Hint:	Verify that the directory containing the `hpc_config` file is writeable by root.
ErrorCode 33:	`INSTALL_CONFIG is set to nfs but NFS_SERVER is set to null.`
ErrorCode 34:	`INSTALL_CONFIG is set to $INSTALL_CONFIG but NFS_SERVER is also set.`
ErrorCode 35:	`An attempt was made to remove a release that is currently active. Please deactivate the release before trying to remove it.`

APPENDIX **B**

Required File and Directory Settings

Sun HPC ClusterTools software requires special settings on several files and filesystems.

HPC Installation Requirement Summary

Files | Directories	Owner	Permissions	Comments
`hpc.conf`	root (or LSF administrator)	644	Must be accessible by all nodes
`hpc_config`	—	644	—
CD-ROM Mount Point	—	—	Must be readable by root and accessible via a common path from all nodes.
HPC *config_dir* (This directory may be removed after installation. However, you should preserve configuration files.)	—	755	Must be r/w by root or shared using the `-o anon = 0` option. (Note: NFS directories may require 777 permissions.)
NFS client mount point (`INSTALL_LOC`)	—	—	Server mount point must exist on all client nodes