



Netra High Availability Suite Foundation Services 2.1 6/03 Quick Start Guide

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054
U.S.A.

Part No: 817-1764-11
September 2004

Copyright 2004 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, docs.sun.com, AnswerBook, AnswerBook2, Java, JMX, Netra, Solaris JumpStart, Solstice DiskSuite, Sun Fire, Javadoc, JDK, Sun4U, Jini, OpenBoot, Sun Workshop, Forte, Sun StorEdge, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc. Adobe is a registered trademark of Adobe Systems, Incorporated.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

Federal Acquisitions: Commercial Software—Government Users Subject to Standard License Terms and Conditions.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2004 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées du système Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux États-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, docs.sun.com, AnswerBook, AnswerBook2, Java, JMX, Netra, Solaris JumpStart, Solstice DiskSuite, Sun Fire, Javadoc, JDK, Sun4U, Jini, OpenBoot, Sun Workshop, Forte, Sun StorEdge, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux États-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux États-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc. Adobe est une marque enregistrée de Adobe Systems, Incorporated.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



040813@9495



Contents

Preface	5
1 Getting Started	9
Overview of Tasks (Task Map)	9
Hardware Requirements	10
Software Requirements	10
2 Setting Up and Connecting the Hardware	11
Setting Up the Netra T1 Servers	11
Setting Up the Terminal Server	11
Setting Up and Configuring the Ethernet Switches	12
Setting Up and Configuring the Installation Server	12
Connecting the Hardware	13
3 Installing the Foundation Services on the Cluster	15
Preparing the Installation Environment	15
▼ To Create a Solaris Distribution on the Installation Server	15
▼ To Prepare the Installation Server	16
Installing and Configuring the <code>nhinstall</code> Tool	17
▼ To Install the <code>nhinstall</code> Tool	17
▼ To Configure the <code>nhinstall</code> Tool	17
Installing the Foundation Services	18
▼ To Start the <code>nhinstall</code> Tool	19
▼ To Boot the Master-Eligible Nodes to Install the Solaris Operating System	19
▼ To Boot the Diskless Node	20

Restarting the ninstall Tool	21
▼ To Restart the ninstall Tool	21
4 Running Administration Tasks on the Cluster	23
Checking the Cluster Nodes	23
▼ To Check the Status of the Cluster Nodes	23
▼ To Check the Network Connection Between Nodes	24
▼ To Check Node Addresses	25
Managing Switchovers and Failovers	26
▼ To Trigger a Switchover	26
▼ To Reboot the Master Node Causing a Failover	27
Configuring a Floating External Address	29
▼ To Configure a Floating External Address	30

Preface

The *Netra High Availability Suite Foundation Services 2.1 6/03 Quick Start Guide* provides a quick and efficient way to install and evaluate the Netra™ High Availability (HA) Suite Foundation Services 2.1 6/03 product. This book provides instructions for installing the Solaris™ operating system and the Foundation Services on a basic cluster.

This book does not describe all of the installation and configuration options available with the Foundation Services. This restriction enables you to have a basic Foundation Services cluster up and running in a short time so that you can evaluate the product.

Who Should Use This Book

This book is for evaluators of the Foundation Services. To install the Foundation Services, you must be familiar with the process of installing the Solaris operating system.

Before You Read This Book

Read the *Netra High Availability Suite Foundation Services 2.1 6/03 Overview* for an overview of the product.

How This Book Is Organized

The book is organized as follows:

[Chapter 1](#) provides an overview of the tasks required to install and manage a basic cluster running the Foundation Services. This chapter also includes the software and hardware requirements for a basic cluster.

[Chapter 2](#) describes how to install and connect the hardware for your cluster.

[Chapter 3](#) describes how to use the `nhinstall` tool to install the Solaris operating system and the Foundation Services on the nodes of your cluster.

[Chapter 4](#) describes some administration tasks to check that the cluster is functioning correctly and to further evaluate the product.

Related Books

You will require some of the following books from the Foundation Services documentation set:

- *Netra High Availability Suite Foundation Services 2.1 6/03 Overview*
- *Netra High Availability Suite Foundation Services 2.1 6/03 Glossary*
- *What's New in Netra High Availability Suite Foundation Services 2.1 6/03*
- *Netra High Availability Suite Foundation Services 2.1 6/03 Quick Start Guide*
- *Netra High Availability Suite Foundation Services 2.1 6/03 Hardware Guide*
- *Netra High Availability Suite Foundation Services 2.1 6/03 Custom Installation Guide*
- *Netra High Availability Suite Foundation Services 2.1 6/03 Cluster Administration Guide*
- *Netra High Availability Suite Foundation Services 2.1 6/03 Troubleshooting Guide*
- *Netra High Availability Suite Foundation Services 2.1 6/03 CMM Programming Guide*
- *Netra High Availability Suite Foundation Services 2.1 6/03 NMA Programming Guide*
- *Netra High Availability Suite Foundation Services 2.1 6/03 Reference Manual*
- *Netra High Availability Suite Foundation Services 2.1 6/03 Standalone CGTP Guide*
- *Netra High Availability Suite Foundation Services 2.1 6/03 Release Notes*
- *Netra High Availability Suite Foundation Services 2.1 6/03 README*

The following books are not part of the Foundation Services documentation set, but these books provide related information:

- *Solaris 8 Advanced Installation Guide*
- *Solaris 9 Installation Guide*

Accessing Sun Documentation Online

The docs.sun.comSM Web site enables you to access Sun technical documentation online. You can browse the docs.sun.com archive or search for a specific book title or subject. The URL is `http://docs.sun.com`.

Ordering Sun Documentation

Sun Microsystems offers select product documentation in print form. For a list of documents and how to order them, see “Buy printed documentation” at `http://docs.sun.com`.

Typographic Conventions

The following table describes the typographic changes that are used in this book.

TABLE P-1 Typographic Conventions

Typeface or Symbol	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name% you have mail.</code>
AaBbCc123	What you type, contrasted with onscreen computer output	<code>machine_name% su</code> <code>password:</code>
<i>AaBbCc123</i>	Command-line placeholder: replace with a real name or value	The command to remove a file is <code>rm filename</code> .

TABLE P-1 Typographic Conventions (Continued)

Typeface or Symbol	Meaning	Example
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. Do <i>not</i> save the file. (Emphasis sometimes appears in bold online.)

Shell Prompts in Command Examples

The following table shows the default system prompt and superuser prompt for the C shell, Bourne shell, and Korn shell.

TABLE P-2 Shell Prompts

Shell	Prompt
C shell prompt	machine_name%
C shell superuser prompt	machine_name#
Bourne shell and Korn shell prompt	\$
Bourne shell and Korn shell superuser prompt	#

Getting Started

The Foundation Services run on a *cluster of nodes* that are interconnected by a network. This book describes how to use the `nhinstall` tool to install and configure the Solaris operating system and the Foundation Services on the nodes of the cluster. To evaluate the Foundation Services product, install the software on a small cluster of two *master-eligible nodes* and a *diskless node*. For more information on these types of nodes and definitions of other terms used in this document, see the *Netra High Availability Suite Foundation Services 2.1 6/03 Glossary*.

To get started, see the following sections:

- “Overview of Tasks (Task Map)” on page 9
- “Hardware Requirements” on page 10
- “Software Requirements” on page 10

Overview of Tasks (Task Map)

The following table summarizes the tasks for installing and checking a Foundation Services cluster. Perform the tasks in sequential order.

	Task	For Instructions
1.	Check that you have the necessary hardware.	“Hardware Requirements” on page 10
2.	Check that you have the necessary software.	“Software Requirements” on page 10
3.	Install and connect the hardware.	Chapter 2, “Setting Up and Connecting the Hardware”

Task	For Instructions
4. Create a Solaris distribution on the installation server and prepare the server for installation.	“Preparing the Installation Environment” on page 15
5. Install and configure the <code>nhinstall</code> tool on the installation server.	“Installing and Configuring the <code>nhinstall</code> Tool” on page 17
6. Start the <code>nhinstall</code> tool and install the software on the cluster nodes.	“Installing the Foundation Services” on page 18
7. Check the nodes of your cluster.	Chapter 4, “Running Administration Tasks on the Cluster”

Hardware Requirements

The following hardware configuration represents a basic cluster with two master-eligible nodes and one diskless node:

- Two Netra T1 servers configured as master-eligible nodes
- One Netra T1 servers configured as a diskless node
- Two Ethernet switches
- A *terminal server* to manage the consoles
- An *installation server* to install the software on the nodes

For information about setting up and connecting the hardware, see [Chapter 2](#).

For other example hardware configurations, see the *Netra High Availability Suite Foundation Services 2.1 6/03 Hardware Guide*.

Software Requirements

You require the following software to install the Foundation Services on the nodes of your cluster:

- A Solaris software distribution. See the Netra High Availability Suite Foundation Services 2.1 6/03 Release Notes for the supported software versions.
- The Foundation Services software distribution

For the procedures to install the software on the nodes of your cluster, see [Chapter 3](#).

Setting Up and Connecting the Hardware

When you have the necessary hardware and software, you are ready to set up the hardware for your cluster.

For instructions on setting up and configuring the cluster hardware, see the following sections. Perform these tasks in sequential order.

- “Setting Up the Netra T1 Servers” on page 11
- “Setting Up the Terminal Server” on page 11
- “Setting Up and Configuring the Ethernet Switches” on page 12
- “Setting Up and Configuring the Installation Server” on page 12
- “Connecting the Hardware” on page 13

Setting Up the Netra T1 Servers

To install the Netra T1 servers, see the documentation provided with the hardware or go to the following site:

<http://www.sun.com/products-n-solutions/hardware/docs/Servers>

Setting Up the Terminal Server

To install the terminal server, see the documentation provided with your terminal server.

Setting Up and Configuring the Ethernet Switches

A Foundation Services cluster must have a redundant network, that is, two network interfaces that back each other up. To make a network redundant, a cluster requires two Ethernet switches. The Foundation Services have been validated on Cisco Catalyst 29x0 Desktop Switches. If you use other switches, make sure that the switches support the following:

- Simple Network Management Protocol (SNMP)
- Management information base (MIB) RFC 1213 and RFC 1493
- Configuration that permits disabling of the Spanning Tree Protocol (STP)

Setting Up and Configuring the Installation Server

The installation server enables you to install the Solaris operating system and the Foundation Services on the cluster by using the Solaris JumpStart™ software.

The installation server requires the following:

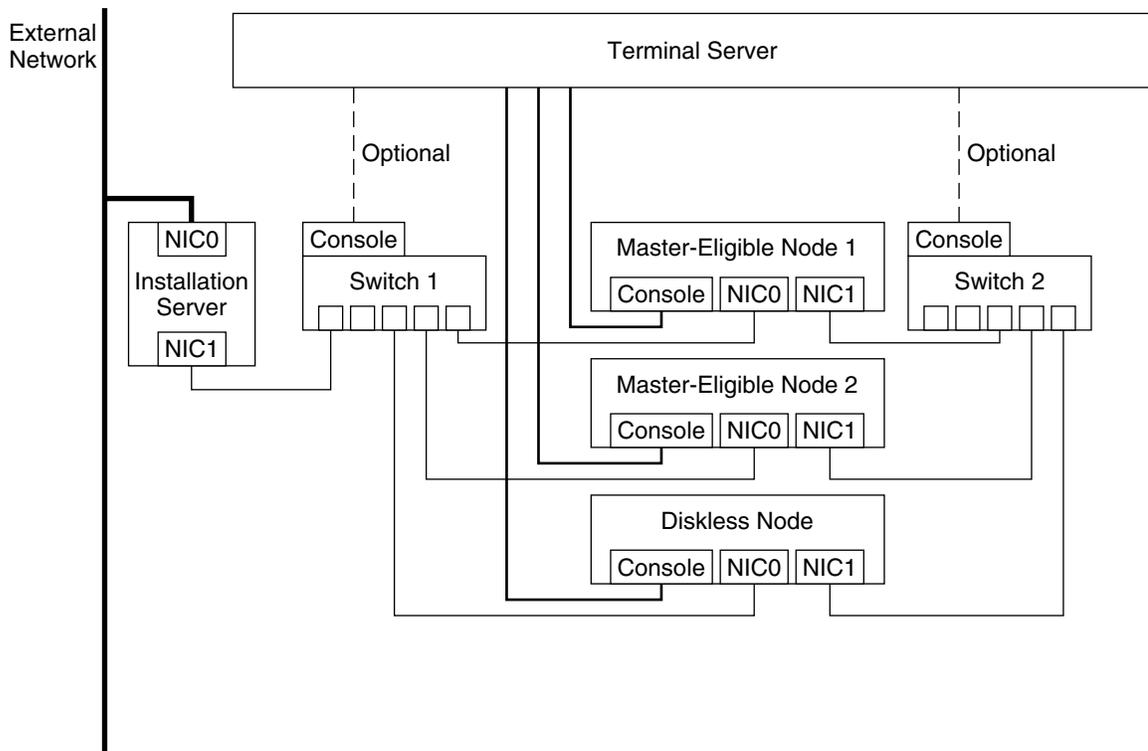
Hardware requirements	UltraSPARC® platform Two network devices. One network device is used to connect the installation server to an external network. The other network device is used to connect the installation server to the cluster network
Operating system	Solaris. See the Netra High Availability Suite Foundation Services 2.1 6/03 Release Notes for information about the supported software versions
Software requirements	Perl Version 5, which is available in the Developer Solaris Software Group The Foundation Services software distribution
Disk capacity	Minimum 1.5 Gbytes for a Solaris software distribution This size does not include applications that you might want to run on the cluster

Free space

Minimum 1.5 Gbytes after the Solaris operating system has been installed

Connecting the Hardware

The following diagram illustrates a cluster of three nodes: two master-eligible nodes and one diskless node.



NIC0 = Interface to the first network card.

NIC1 = Interface to the second network card.

FIGURE 2-1 Connecting the Hardware

The nodes of your cluster are connected to each other through switches. You can connect the console of each node to the terminal server to provide access to the console of the node. Connect the installation server to the cluster network through a switch. To enable console access, connect the serial port of the nodes to the terminal server.

The cluster nodes must not be on the same physical wire as the nodes of another cluster. When a diskless node boots, it sends a broadcast message to find the master node. If two clusters share the same wire, the diskless node could receive messages from the wrong master node.

Installing the Foundation Services on the Cluster

The `nhinstall` tool enables you to install and configure the Foundation Services on cluster nodes. This tool must be installed on the installation server, which must be connected to your cluster. For details on how to connect cluster nodes and the installation server, see [Chapter 2](#).

To install the software on the cluster by using the `nhinstall` tool, see the following sections. Perform the tasks in sequential order.

- [“Preparing the Installation Environment” on page 15](#)
- [“Installing the Foundation Services” on page 18](#)
- [“Restarting the `nhinstall` Tool” on page 21](#)

Preparing the Installation Environment

Before installing the `nhinstall` tool on the installation server, create a Solaris distribution on the installation server. You must also prepare the installation server to install the Solaris operating system and the Foundation Services on nodes of the cluster.

▼ To Create a Solaris Distribution on the Installation Server

The Solaris distribution is used to install the Solaris operating system on the nodes of the cluster.

1. **Check that at least 1.5 Gbytes of free disk space exists on the installation server.**
2. **Log in to the installation server as superuser.**

3. Create a directory for the Solaris distribution.

4. Change to the directory where the `setup_install_server` tool is located.

```
# cd Solaris-Dir/Solaris_x/Tools
```

- *Solaris-Dir* is the directory containing the Solaris installation software. This could be a CD-ROM or an NFS-shared directory.
- The *x* in *Solaris_x* depends on the Solaris version you plan to install. For the Solaris 9 operating system it is *Solaris_9*.

5. Run the `setup_install_server` command.

```
# ./setup_install_server Solaris-Distribution
```

where *Solaris-Distribution* is the directory you created in [Step 3](#).

For more information on the `setup_install_server` command, see the appropriate documentation:

- *Solaris 8 Advanced Installation Guide* and the `setup_install_server(1M)` man page
- *Solaris 9 Installation Guide* and the `setup_install_server(1M)` man page

▼ To Prepare the Installation Server

1. Configure the installation server as described in [“Setting Up and Configuring the Installation Server” on page 12](#).

2. Use the `pkginfo` command to confirm that the `SUNWp15u`, `SUNWp15p`, and `SUNWp15m` Perl 5.0 packages are installed on the installation server.

3. Delete any entries for your cluster nodes in the following files:

- `/etc/hosts`
- `/etc/ethers`, if it exists
- `/etc/bootparams`, if it exists

The `nhinstall` tool automatically modifies these files on the installation server to add entries for the cluster nodes. If the `nhinstall` tool finds that the entries in these files differ from their content at initial install time, the installation stops and the tool displays an error.

4. Disable the installation server as a router by creating a `/etc/notrouter` file.

```
# touch /etc/notrouter
```

If a system running the Solaris operating system has two network interfaces, the system is configured as a router by default. However, for security reasons, a Foundation Services cluster network must not be routed.

5. In the `/etc/nsswitch.conf` file on the installation server, place `files` before `nis` in the `hosts`, `ethers`, and `bootparams` entries.

```
hosts: files nis
ethers: files nis
bootparams: files nis
```

6. From the installation server or the system from which you are connected to the installation server, open a terminal window to connect to the console of each node.

Installing and Configuring the `nhinstall` Tool

▼ To Install the `nhinstall` Tool

Install the package containing the `nhinstall` tool on the installation server.

1. Log in to the installation server as superuser.
2. Install the `nhinstall` package, `SUNWnhins`, on the installation server.

```
# pkgadd -d /software-distribution-dir/NetraHAS2.1/Packages/ SUNWnhins
```

where `software-distribution-dir` is the directory that contains the Foundation Services packages.

3. To access the man pages on the installation server, install the `SUNWnhman` package.

```
# pkgadd -d /software-distribution-dir/NetraHAS2.1/Packages/ SUNWnhman
```

The man pages are installed in the `/opt/SUNWcgha/man` directory.

▼ To Configure the `nhinstall` Tool

1. Log in to the installation server as superuser.
2. Copy the configuration files to a local directory on the installation server.

```
# mkdir config-file-directory
# export NHOME=/opt/SUNWcgha/config.standard
# cd config-file-directory
# cp $NHOME/env_installation.conf.template env_installation.conf
# cp $NHOME/cluster_definition.conf.template cluster_definition.conf
```

The templates for the configuration files are contained in the `/opt/SUNWcgha/config.standard` directory with `.template` extensions. The configuration files must be in the same local directory on the installation server.

3. To define the installation environment, configure parameters in the `env_installation.conf` file, for example:

```
SERVER_INTERFACE=hme1
AUTO_REBOOT=YES
SOLARIS_INSTALL=ALL
NHAS2_PRODUCT_DIR=/downloads/NetraHASuite
SOLARIS_DIR=/export/home/s9hw13
```

Note – Make sure that the directories defined for the parameters `NHAS2_PRODUCT_DIR` and `SOLARIS_DIR` are shared through the Network File System (NFS).

For more information, see the `env_installation.conf(4)` man page.

4. To define the cluster nodes, configure parameters in the `cluster_definition.conf` file such as:

```
MEN_INTERFACES=hme0 hme1
NMEN_INTERFACES=hme0 hme1
NODE=10 08:00:20:f9:c5:54 - node10
NODE=20 08:00:20:f9:a8:12 - node20
NODE=30 - - node30
```

- `node10` and `node20` are the master-eligible nodes.
- `node30` is the diskless node.
- For the master-eligible nodes, make sure that you specify the Ethernet address of the network interface that is connected to the same switch as the installation server. Usually, the *NIC0*, that is, the `hme0`, and the installation server are connected to the first switch.

For details on additional parameters, including the parameters for defining external network access for the nodes, see the `cluster_definition.conf(4)` man page.

Installing the Foundation Services

After you have prepared the installation environment and installed and configured the `ninstall` tool on the installation server, you are ready to install the software on the nodes of the cluster. The entire installation takes approximately one and a half hours.

▼ To Start the `nhinstall` Tool

1. Log in to the installation server as superuser.
2. Start the `nhinstall` tool on the installation server.

```
# cd /opt/SUNWcgha/sbin
# ./nhinstall -r config-file-directory -l logfile
```

For details, see the `nhinstall(1M)` man page.

The `nhinstall` tool prepares the environment by using the Solaris JumpStart utility.

Note – If the Solaris JumpStart procedure stops, the utility might not restart. The cause can be the presence of the `/tmp/.install_client.lock` file, which prevents two Solaris JumpStart commands from being executed simultaneously. To restart the Solaris JumpStart, make sure that there are no other Solaris JumpStart processes running, and then delete this lock file.

After the Solaris JumpStart environment has been prepared, the following instructions are displayed by the `nhinstall` tool:

```
On the console of each master-eligible node:
1. Get the ok prompt.
2. At the prompt, type:
    ok> boot net - install
Answer 'y' after the command has been entered on both
master-eligible nodes:
```

Follow these instructions by performing the procedures in the remainder of this section.

▼ To Boot the Master-Eligible Nodes to Install the Solaris Operating System

1. Go to the console window of the first master-eligible node.

The first master eligible-node is the first node defined in the `cluster_definition.conf` file with the `NODE` definition. In the example used in this document, the first master-eligible node is `node10`.

2. Get the `ok` prompt.

- a. To get the `telnet` prompt, press `Control-]`.

```
# Control-]
```

- b. To get the `ok` prompt, type `send brk` at the `telnet` prompt.

```
telnet> send brk
Type 'go' to resume
```

```
ok>
```

3. Set the boot variables by using the `setenv` command at the `ok` prompt.

- a. Configure the processors to use local Ethernet (MAC) addresses.

```
ok> setenv local-mac-address? true
```

- b. Configure the processors to retry booting when automatic boot fails.

```
ok> setenv auto-boot-retry? true
```

4. Repeat Step 1 through Step 3 for the second master-eligible node.

The second master-eligible node is the second node defined with the `NODE` parameter in the `cluster_definition.conf` file. In the example used in this document, the second master-eligible node is `node20`.

5. Start the installation of the Solaris operating system on the first master-eligible node.

```
ok> boot net - install
```

Where `net` is the device alias.

6. At the same time, start the installation of the Solaris operating system on the second master-eligible node.

```
ok> boot net - install
```

7. After you type the command to boot the master-eligible nodes, type `y` at the `nhinstall` prompt in the installation server console.

```
Answer 'y' after the command has been entered on  
both master-eligible nodes: y
```

The `nhinstall` tool checks the version of the Solaris operating system that has been installed. The `nhinstall` tool also checks whether the Foundation Services are already installed on the master-eligible nodes. If the Foundation Services is already installed, the `nhinstall` tool exits and displays an error message. The `nhinstall` tool then installs the Foundation Services on the nodes of the cluster.

▼ To Boot the Diskless Node

The following instructions are displayed by the `nhinstall` tool in the console of the installation server:

```
Setup the eeprom boot parameters on your diskless nodes:
```

```
At the prompt, type:
```

```
ok> setenv local-mac-address? true
```

```
ok> setenv auto-boot-retry? true
```

```
ok> setenv boot-device net:dhcp,,,,,5 net2:dhcp,,,,,5
```

```
ok> setenv diag-switch? false
```

```
You can now boot your diskless nodes.
```

Configure the boot parameters for the diskless node. If you are installing more than one diskless node, repeat the following steps for each diskless node.

1. Get the `ok` prompt on the diskless node.
2. Execute these commands at the `ok` prompt on the diskless node.

```
ok> setenv local-mac-address? true
ok> setenv auto-boot-retry? true
ok> setenv boot-device net:dhcp,,,,,5 net2:dhcp,,,,,5
ok> setenv diag-switch? false
```

3. Boot the diskless node.

```
ok> reset
```

While you are booting the node, the following information is displayed by the `nhinstall` tool:

```
17:32:21 Waiting for the disks of the master-eligible nodes to synchronize.
17:32:21 - Partitions to be replicated:
17:32:22 . md/rdisk/d4 (2099 MB)
17:32:22 . md/rdisk/d3 (2099 MB)
0%...10%...20%...30%...40%...50%...60%...70%...80%...90%...100%
17:44:18 Your cluster is ready.
```

When the cluster is first started, there is a period during which the *master node* and *vice-master node* synchronize their disks. This output describes the status of the initial synchronization of the disks of master node and the vice-master node. After the diskless node is rebooted and after the synchronization is complete, your cluster is ready.

Restarting the `nhinstall` Tool

Most warnings displayed by the `nhinstall` tool require no action. When the `nhinstall` tool is launched, the tool parses the configuration files and stops if the tool encounters errors in the files. After you correct the errors, you can continue the installation from the point at which the error occurred.

▼ To Restart the `nhinstall` Tool

- To restart the `nhinstall` tool, use the `nhinstall` command with the same options that you used to start the tool.

```
# cd /opt/SUNWcgha/sbin
# ./nhinstall -r config_file_directory -l logfile
```

If you modify the configuration files to correct an error, the `nhinstall` tool displays a warning that the configuration has changed. The `nhinstall` tool then prompts you to reset the installation. If you do not want to reset the installation, type `no` at the prompt.

Running Administration Tasks on the Cluster

After you have installed the software on the cluster, run some administration tasks to check that the cluster is functioning correctly and to further evaluate the product.

To check your cluster, perform the following administration tasks:

- “Checking the Cluster Nodes” on page 23
- “Managing Switchovers and Failovers” on page 26
- “Configuring a Floating External Address” on page 29

Checking the Cluster Nodes

The Foundation Services product is delivered with tools to check different aspects of a cluster, including the status of cluster nodes, the network connection between nodes, and the IP addresses of nodes.

▼ To Check the Status of the Cluster Nodes

To check the nodes of your cluster, use the `nhcmmstat` command.

1. **Log in to a master-eligible node as superuser.**
2. **Check the nodes by using the `nhcmmstat` command.**

```
# /opt/SUNWcgha/sbin/nhcmmstat -c all
Executed Command: all
-----
node_id      = 10    [This is the current node]
domain_id   = 250
name        = MEN-C250-N10
```

```

role           = MASTER
qualified      = YES
synchro.      = READY
frozen        = NO
excluded      = NO
eligible      = YES
incarn.       = 1038420771 (27/11/2002 - 19:12:51)
swload_id     = 1
CGTP @       = 10.240.3.10
-----

```

```

node_id       = 30
domain_id    = 250
name         = node30
role        = IN
qualified    = YES
synchro.    = READY
frozen      = NO
excluded    = NO
eligible    = NO
incarn.     = 1038422116 (27/11/2002 - 19:35:16)
swload_id   = 1
CGTP @     = 10.240.3.30
-----

```

```

node_id       = 20
domain_id    = 2540
name         = MEN-C250-N20
role        = VICE-MASTER
qualified    = YES
synchro.    = READY
frozen      = NO
excluded    = NO
eligible    = YES
incarn.     = 1038420945 (27/11/2002 - 19:15:45)
swload_id   = 1
CGTP @     = 10.240.3.20
-----

```

In the preceding example, the output from the `nhcmmstat` command displays information about all the peer nodes in the console window. This information includes the role of each node. The peer nodes must include the master and vice-master nodes.

For more information on `nhcmmstat`, see the `nhcmmstat(1M)` man page.

▼ To Check the Network Connection Between Nodes

To check that the cluster network is functioning correctly, use the `nhadm` command.

1. **Log in to a peer node as superuser.**
2. **Verify that the nodes in the cluster are communicating through a network.**

```
# /opt/SUNWcgha/sbin/nhadm check
```

If any peer node is not accessible from any other peer node, the `nhadm` command displays an error message in the console window.

For more information, see the *Netra High Availability Suite Foundation Services 2.1 6/03 Cluster Administration Guide*.

▼ To Check Node Addresses

Each node has an IP address assigned to the `NIC0`, `NIC1`, and `cgtp0` network interfaces. To identify and ping each network interface of a node, follow this procedure.

1. Log in to the node that you want to examine.
2. Run the `ifconfig` command.

```
# ifconfig -a
```

The `ifconfig` command displays configuration information about the network interfaces to the console window. Sample output for the `ifconfig` command on a peer node is as follows:

```
hme0: flags=1004843<UP,BROADCAST,RUNNING,MULTICAST,DHCP,IPv4> mtu 1500 \
index 1
    inet 10.250.1.30 netmask ffffffff broadcast 10.250.1.255
    ether 8:0:20:f9:b4:b0
lo0: flags=1000849<UP,LOOPBACK,RUNNING,MULTICAST,IPv4> mtu 8232 index 2
    inet 127.0.0.1 netmask ff000000
hme1: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 3
    inet 10.250.2.30 netmask ffffffff broadcast 10.250.2.255
    ether 8:0:20:f9:b4:b1
cgtp0: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 4
    inet 10.250.3.30 netmask ffffffff broadcast 10.250.3.255
    ether 0:0:0:0:0:0
```

Each peer node has at least three network interfaces configured. If a node has external access configured or if the node is the master, more network interfaces are displayed by the `ifconfig` command.

3. Retrieve the cluster ID, that is, the `domainid`, by using the output from the `ifconfig` command.
The `domainid` in this example is 250.
4. Retrieve the node ID, that is, the `nodeid`, by using the output from the `ifconfig` command.
The `nodeid` in this example is 30.
5. Retrieve the network interface names and corresponding IP addresses by using the output from the `ifconfig` command.

The network interfaces *NIC0* and *NIC1* in this example are the physical interfaces *hme0* and *hme1*, respectively. The third interface is the virtual physical interface, *cgtp0*.

The IP addresses for the three network interfaces in this example are as follows:

```
hme0    10.250.1.30
hme1    10.250.2.30
cgtp0    10.250.3.30
```

The Ethernet addresses for *NIC0* and *NIC1* in this example are as follows:

```
hme0    8:0:20:f9:b4:b0
hme1    8:0:20:f9:b4:b1
```

6. Log in to another peer node as superuser.
7. Ping each network interface address of the node 30.

```
# ping 10.250.1.30
# ping 10.250.2.30
# ping 10.250.3.30
```

Managing Switchovers and Failovers

You can trigger a *switchover* to swap the master and vice-master roles of the master-eligible nodes. A switchover is useful when you plan to take the master node down for maintenance. To trigger a switchover, see [“To Trigger a Switchover” on page 26](#).

However, if there is a problem on the master node, the master role fails over automatically to the vice-master node. In this case, the master and vice-master roles are also swapped, but because the cause is an unplanned problem, the swap is called a *failover*. To cause a failover, see [“To Reboot the Master Node Causing a Failover” on page 27](#).

▼ To Trigger a Switchover

1. Log in to a peer node as superuser.
2. Identify the master node.

```
# /opt/SUNWcgha/sbin/nhcmstat -c all
```

The `nhcmmstat` command prints information on each peer node to the console window.

3. Log in to the master node as superuser.

4. Trigger a switchover.

```
# /opt/SUNWcgha/sbin/nhcmmstat -c so
```

If there is a vice-master node qualified to become master in the cluster, this node is elected master. The old master node becomes the vice-master node. If there is no potential master, `nhcmmstat` does not perform a switchover.

5. After the switchover is complete, verify that the roles of the master and vice-master nodes have been switched.

```
# /opt/SUNWcgha/sbin/nhcmmstat -c vice
```

If the switchover is successful, the current node is the vice master. This command also verifies that the current node is synchronized with the new master node.

6. Verify the cluster configuration.

```
# /opt/SUNWcgha/sbin/nhadm check
```

For more information on `nhcmmstat`, see the `nhcmmstat(1M)` man page.

▼ To Reboot the Master Node Causing a Failover

If you reboot the master node, you trigger a failover.

1. Log in to a peer node as superuser.

2. Run the `nhcmmstat` command to identify the master node.

```
# /opt/SUNWcgha/sbin/nhcmmstat -c all
```

3. Log in to the master node as superuser.

4. Shut down the master node.

```
# init 5
```

The vice-master node becomes the master. Because one of the two master-eligible nodes in the cluster is shut down, you lose the redundancy of the cluster. To recover redundancy, restart the stopped node.

5. Log in to a peer node as superuser.

6. Verify that the vice-master node became the master node when the old master node was shut down.

```

# /opt/SUNWcgha/sbin/nhcmstat -c master
Executed Command: master
-----
node_id      = 20    [This is the current node]
domain_id    = 250
name         = MEN-C250-N20
role         = MASTER
qualified    = YES
synchro.     = NEEDED !!!
frozen       = NO
excluded     = NO
eligible     = YES
incarn.      = 1038481013 (28/11/2002 - 11:56:53)
swload_id    = 1
CGTP @      = 10.250.3.20
-----

```

The output shows that the vice-master node is now the master node. In addition, the new master node displays a requirement for synchronizing its disk with the disk of the old master node.

7. Restart the old master node, which you shut down in [Step 4](#).

```
lom> poweron
```

This node now automatically becomes the vice-master node.

8. Run the `nhcmstat` command to verify that the current node is the vice-master node.

```

# /opt/SUNWcgha/sbin/nhcmstat -c all
Executed Command: all
-----
node_id      = 30
domain_id    = 250
name         = node30
role         = IN
qualified    = YES
synchro.     = READY
frozen       = NO
excluded     = NO
eligible     = NO
incarn.      = 1038422116 (27/11/2002 - 19:35:16)
swload_id    = 1
CGTP @      = 10.250.3.30
-----
-----
node_id      = 20
domain_id    = 250
name         = MEN-C250-N20
role         = MASTER
qualified    = YES
synchro.     = READY
frozen       = NO
excluded     = NO

```

```

eligible      = YES
incarn.      = 1038481013 (28/11/2002 - 11:56:53)
swload_id    = 1
CGTP @      = 10.250.3.20
-----
node_id      = 10    [This is the current node]
domain_id    = 250
name         = MEN-C250-N10
role         = VICE-MASTER
qualified    = YES
synchro.     = READY
frozen       = NO
excluded     = NO
eligible     = YES
incarn.      = 1038481383 (28/11/2002 - 12:03:03)
swload_id    = 1
CGTP @      = 10.250.3.10
-----

```

9. Log in to the new vice-master node as superuser.

10. Verify that the node has started correctly.

```
# /opt/SUNWcgha/sbin/nhadm check
```

For more information on the tests run by `nhadm check`, see the `nhadm(1M)` man page.

Configuring a Floating External Address

A *floating external address* is a logical address assigned to an interface that is used to connect the master node to an external network. The Node State Manager (NSM) uses the Cluster Membership Manager (CMM) notifications to determine when a node takes on or loses the master role. When notified that a node has become the master node, the NSM executes a script to configure a floating external address on one of the node's external interfaces. When notified that a node has lost the master role, the NSM executes a script to unconfigure the floating external address.

The NSM can be installed when you first install the software on the cluster or after you have completed the installation process and have a running cluster. The following procedure describes how to install the NSM on a running cluster.

▼ To Configure a Floating External Address

1. Log in to the master node as superuser.
2. Create a file named `not_configured` in the `/etc/opt/SUNWcgha` directory.

```
# touch /etc/opt/SUNWcgha/not_configured
```

If the node is rebooted during this procedure, the node does not start the Foundation Services.

3. Reboot the master node.
4. Install the NSM packages, `SUNWnhnsa` and `SUNWnhnsb` on the master node.

```
# pkgadd -d /software-distribution-dir/NetraHAS2.1/Packages/ \
SUNWnhnsa SUNWnhnsb
```

where `software-distribution-dir` is the directory that contains the Foundation Services packages.

5. Edit the `/etc/opt/SUNWcgha/nhfs.conf` file to define the NSM parameters.

Example entries to configure the NSM are as follows:

```
NSM.Exec.Masterdir=/opt/SUNWcgha/actions/master
NSM.Exec.Vicemasterdir=/opt/SUNWcgha/actions/vicemaster
NSM.Log.Masterdir=/var/adm/log/SUNWcgha/nsm_master
NSM.Log.Vicemasterdir=/var/adm/log/SUNWcgha/nsm_vicemaster
NSM.External.Master.Address=192.168.12.39
NSM.External.Master.Nic=hme0:101
```

The floating external address is attributed to the master role. This address is assigned to the logical interface `hme0:101` and has the external IP address `192.168.12.39`. If the master node switches over to become the vice-master node, the interface `hme0:101` is enabled on the new master node. For more details on these parameters, see the `nhfs.conf(4)` man page.

6. Log in to the vice-master node.
7. Repeat [Step 2](#) through [Step 5](#) on the vice-master node.
8. On both the master node and the vice-master node, delete the `/etc/opt/SUNWcgha/not_configured` file.
9. Reboot both the master node and the vice-master node.
10. Log in to the master node.
11. Run the `ifconfig` command on the master node.

```
# ifconfig -a
```

```

lo0: flags=1000849<UP,LOOPBACK,RUNNING,MULTICAST,IPv4> mtu 8232 \
index 1
    inet 127.0.0.1 netmask ff000000
hme0: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 2
    inet 10.250.1.20 netmask ffffffff broadcast 10.250.1.255
    ether 8:0:20:f9:b3:6a
hme0:1: flags=1040843<UP,BROADCAST,RUNNING,MULTICAST,DEPRECATED,IPv4> \
mtu 1500 index 2
    inet 10.250.1.1 netmask ffffffff broadcast 10.250.1.255
hme0:101: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> \
mtu 1500 index 2
    inet 192.168.12.39 netmask ffffffff broadcast 172.36.128.255
hme1: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 3
    inet 10.250.2.20 netmask ffffffff broadcast 10.250.2.255
    ether 8:0:20:f9:b3:6b
hme1:1: flags=1040843<UP,BROADCAST,RUNNING,MULTICAST,DEPRECATED,IPv4> \
mtu 1500 index 3
    inet 10.250.2.1 netmask ffffffff broadcast 10.250.2.255
cgtp0: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 4
    inet 10.250.3.20 netmask ffffffff broadcast 10.250.3.255
    ether 0:0:0:0:0:0
cgtp0:1: flags=1040843<UP,BROADCAST,RUNNING,MULTICAST,DEPRECATED,IPv4> \
mtu 1500 index 4
    inet 10.250.3.1 netmask ffffffff broadcast 10.250.3.255

```

In this output, you can see the entry for the hme0:101 interface with the floating external address 192.168.12.39.

12. Run the ifconfig command on the vice-master node.

```

# ifconfig -a
lo0: flags=1000849<UP,LOOPBACK,RUNNING,MULTICAST,IPv4> mtu 8232 \
index 1
    inet 127.0.0.1 netmask ff000000
hme0: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 \
index 2
    inet 10.250.1.10 netmask ffffffff broadcast 10.250.1.255
    ether 8:0:20:da:8f:c6
hme1: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 \
index 3
    inet 10.250.2.10 netmask ffffffff broadcast 10.250.2.255
    ether 8:0:20:da:8f:c7
cgtp0: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 \
index 4
    inet 10.250.3.10 netmask ffffffff broadcast 10.250.3.255
    ether 0:0:0:0:0:0

```

In this output, there is no entry for the external network interface, hme0:101. Therefore, there is no floating external address on the vice-master node.

13. Trigger a switchover.

```

# /opt/SUNWcgha/sbin/nhcmstat -c so

```

14. Run the `ifconfig` command on the new master node.

```
# ifconfig -a
lo0: flags=1000849<UP,LOOPBACK,RUNNING,MULTICAST,IPv> mtu 8232 index 1
    inet 127.0.0.1 netmask ff000000
hme0: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 \
index 2
    inet 10.250.1.10 netmask ffffffff broadcast 10.250.1.255
    ether 8:0:20:da:8f:c6
hme0:1: flags=1040842<BROADCAST,RUNNING,MULTICAST,DEPRECATED,IPv4> mtu
1500 index 2
    inet 10.250.1.1 netmask ffffffff broadcast 10.250.1.255
hme0:101: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> \
mtu 1500 index 2
    inet 192.168.12.39 netmask ffffffff broadcast 172.36.128.255
hme1: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index
3
    inet 10.250.2.10 netmask ffffffff broadcast 10.250.2.255
    ether 8:0:20:da:8f:c7
hme1:1: flags=1040842<BROADCAST,RUNNING,MULTICAST,DEPRECATED,IPv4> mtu
1500 index 3
    inet 10.250.2.1 netmask ffffffff broadcast 10.250.2.255
cgtp0: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index
4
    inet 10.250.3.10 netmask ffffffff broadcast 10.250.3.255
    ether 0:0:0:0:0:0
cgtp0:1: flags=1040842<BROADCAST,RUNNING,MULTICAST,DEPRECATED,IPv4> mtu
1500 index 4
    inet 10.250.3.1 netmask ffffffff broadcast 10.250.3.255
```

In this output, the new master node has the `hme0:101` interface with the floating external address `192.168.12.39`.

15. From a remote system, ping the master node floating address.

```
% ping -s 192.168.12.39
```

Where `192.168.12.39` is the floating address assigned to the master node.