



Netra High Availability Suite Foundation Services 2.1 6/03 Troubleshooting Guide

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

Part No: 817-1769-11
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, Javadoc, Sun Fire, 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 Etats-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, Javadoc, Sun Fire, 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 Etats-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 Etats-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.



040830@9495



Contents

Preface 7

Part I Introduction to Troubleshooting 11

1 Using the Troubleshooting Guide 13

When to Use This Book 13

How to Use This Book 14

2 Using Administration Tools and Configuration Files to Troubleshoot 17

Part II Troubleshooting at Installation and Startup 19

3 Recovering From Installation Problems 21

Recovering From General Installation Problems 21

 Incorrect Software Is Installed on the Cluster Nodes 21

Recovering From `nhinstall` Problems 22

 The `nhinstall` Tool Stops During Installation 22

 Solaris JumpStart Installation Fails During `nhinstall` Installation 23

4 Recovering From Startup Problems on Master-Eligible Nodes 25

A Master-Eligible Node Does Not Boot 25

 ▼ To Investigate Why the Foundation Services Do Not Start on a Master-Eligible Node 26

A Master Node Is Not Elected at Startup 27

▼ To Investigate Why a Master Node Is Not Elected at Startup	27
Two Master Nodes Are Elected at Startup	29
▼ To Investigate Split Brain on Clusters With a Direct Link	29
▼ To Investigate Split Brain on Clusters Without a Direct Link	30
The Vice-Master Node Remains Unsynchronized After Startup	31
▼ To Investigate Why the Vice-Master Node Remains Unsynchronized After Startup	31
A Monitored Daemon Fails Causing a Master-Eligible Node to Reboot at Startup	33
The Node Management Agent on a Master-Eligible Node Exits at Startup	34
▼ To Investigate Why the NMA on a Master-Eligible Node Exits at Startup	34
5 Recovering From Startup Problems on Diskless Nodes and Dataless Nodes	37
A Diskless Node Does Not Boot at Startup	37
▼ To Investigate Why the Solaris Operating System Does Not Start on a Diskless Node	37
▼ To Investigate Why the Foundation Services Do Not Start on a Diskless Node	40
A Diskless Node Does Not Boot After Failover	41
A Dataless Node Does Not Boot at Startup	41
▼ To Investigate Why the Foundation Services Do Not Start on a Dataless Node	42
A Monitored Daemon Fails Causing a Diskless Node or Dataless Node to Reboot at Startup	42

Part III Troubleshooting at Runtime 45

6 Recovering From Failover and Switchover Problems	47
Two Master Nodes Are Elected at Runtime	47
▼ To Investigate Split Brain During Runtime on Clusters Without a Direct Link	47
A Diskless Node Does Not Reboot After Failover	48
▼ To Reboot a Diskless Node After Failover	48
Replication Does Not Resume After Failover or Switchover	49
7 Recovering From Node Reboot at Runtime	51
A Monitored Daemon Fails Causing a Node to Reboot at Runtime	51
▼ To Recover From Daemon Failure	52

8	Cannot Add Nodes to a Running Cluster	55
	Cannot Add a Node to a Running Cluster by Using the <code>nhinstall</code> Tool	55
	▼ To Investigate Why You Cannot Add a Diskless Node to a Running Cluster by Using the <code>nhinstall</code> Tool	55
9	Cannot Collect Statistics by Using the Node Management Agent	59
	An External Client Cannot Communicate With the Node Management Agent	59
	▼ To Investigate Why an External Client Cannot Communicate With an NMA on a Peer Node	59
	NMA Not Restarted After Failure	60
	▼ To Investigate Why the NMA Is Not Restarted	60
	NMA Not Sending SNMP Traps to a Given Target	61
	The <code>switchOver</code> Method Does Not Finish Executing	61
	Cascading Fails	61
	▼ To Examine Why Cascading Fails	62
A	Error Messages	63
	Introduction to Error Messages	63
	Error Messages Written During Installation	64
	Error Messages Written During Manual Installation	64
	Error Messages Written During Installation Using the <code>nhinstall</code> Tool	64
	Error Messages Written During Runtime	65
	Error Messages Written by the Cluster Membership Manager	65
	Error Messages Written by Reliable NFS	67
	Error Messages Written by the Reliable Boot Service	69
	Error Messages Written by the Watchdog Timer	69
	Error Messages Written by the Node Management Agent	70
	Error Messages Written by Command-Line Tools	70
	Error Messages Written by the <code>nhadm</code> Command	71
	Error Messages Written by the <code>nhcrfsadm</code> Command	74
	Index	75

Preface

This book describes how to solve problems when installing, starting up, and maintaining the Netra™ High Availability (HA) Suite Foundation Services 2.1 6/03 product.

Who Should Use This Book

This book is for system administrators who are installing or maintaining the Foundation Services. This book assists administrators with the following troubleshooting tasks:

- Diagnosing and solving problems encountered while installing, configuring, or maintaining the Foundation Services
- Understanding error messages and resolving errors

How This Book Is Organized

This book is divided into parts that contains related chapters. An appendix is also provided.

[Part I](#) introduces the Troubleshooting Guide and the Foundation Services tools and configuration files.

- [Chapter 1](#) explains when and how to use this book. This chapter also points to other books in the documentation set for answers to common questions.
- [Chapter 2](#) describes the suite of tools and configuration files to use to administer and debug your cluster environment.

Part II describes how to diagnose and solve problems when you are installing and starting up the Foundation Services.

- Chapter 3 describes how to diagnose and solve problems during installation of the Solaris™ operating system and the Foundation Services.
- Chapter 4 describes what to do if the master-eligible nodes do not startup correctly after a seemingly successful installation.
- Chapter 5 describes what to do when you have started up the master-eligible nodes but are unable to startup the diskless nodes or dataless nodes.

Part III describes how to diagnose and solve problems when your cluster is running.

- Chapter 6 describes how to recover from problems associated with failover and switchover.
- Chapter 7 describes how to solve problems that cause a node to reboot during runtime.
- Chapter 8 describes what to do if you cannot add a node to your cluster.
- Chapter 9 describes error scenarios that can occur when you are using the NMA.

Appendix A lists some of the error messages that are produced by the Foundation Services and describes how to resolve the errors.

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*

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> Password:
<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	#

Introduction to Troubleshooting

For an introduction to this book and the Foundation Services tools and configuration files, see the following chapters:

- [Chapter 1](#) explains when and how to use this book. This chapter also points to other books in the Foundation Services documentation set for answers to common questions.
- [Chapter 2](#) describes the suite of tools and configuration files to use to maintain and debug your cluster environment.

Using the Troubleshooting Guide

For information about when and how to use this book, and for pointers to other books in the Foundation Services documentation set, see the following sections:

- “When to Use This Book” on page 13
- “How to Use This Book” on page 14

When to Use This Book

This book provides a defined approach to solving problems that can occur when you are installing, starting up, or using the Foundation Services. This book does not explain how to install or maintain your cluster. Nor does this book explain the tools, configuration files, and APIs provided by the Foundation Services. These subjects are covered in other books in the documentation set.

The following table lists some common questions and answers about the Foundation Services.

TABLE 1-1 Common Questions

Question	For Information, see
What are the Foundation Services?	<i>Netra High Availability Suite Foundation Services 2.1 6/03 Overview</i>
What’s new in this release?	<i>What’s New in Netra High Availability Suite Foundation Services 2.1 6/03</i>

TABLE 1-1 Common Questions (Continued)

Question	For Information, see
What tools and configuration files are available?	<i>Netra High Availability Suite Foundation Services 2.1 6/03 Reference Manual</i> for the man pages that can be accessed by using the <code>man(1)</code> command
How can the Foundation Services be evaluated easily?	<i>Netra High Availability Suite Foundation Services 2.1 6/03 Quick Start Guide</i>
How is the hardware configured?	<i>Netra High Availability Suite Foundation Services 2.1 6/03 Hardware Guide</i>
Which installation method should you use?	<i>Netra High Availability Suite Foundation Services 2.1 6/03 Overview</i>
How are the Foundation Services installed?	<i>Netra High Availability Suite Foundation Services 2.1 6/03 Custom Installation Guide</i>
How is a cluster maintained, and how is hardware added or replaced?	<i>Netra High Availability Suite Foundation Services 2.1 6/03 Cluster Administration Guide</i>
How are applications that use the Foundation Services developed?	<i>Netra High Availability Suite Foundation Services 2.1 6/03 CMM Programming Guide</i> or the <i>Netra High Availability Suite Foundation Services 2.1 6/03 NMA Programming Guide</i>
What does a term or a name mean?	<i>Netra High Availability Suite Foundation Services 2.1 6/03 Glossary</i>

How to Use This Book

The subsequent chapters in this book provide detailed information about how to solve specific installation, startup, or runtime problems. Each chapter covers a different problem area. The following table shows where to find troubleshooting information for a particular troubleshooting area.

TABLE 1-2 Where to Find Troubleshooting Information

Problem Area	For Information, see
Installation and startup	Part II
Installation of the Solaris operating system and the Foundation Services	Chapter 3
Startup of master-eligible nodes	Chapter 4

TABLE 1-2 Where to Find Troubleshooting Information (Continued)

Problem Area	For Information, see
Startup of diskless nodes or dataless nodes	Chapter 5
Runtime	Part III
Failover and switchover during runtime	Chapter 6
Node reboot during runtime	Chapter 7
Addition of nodes to your cluster during runtime	Chapter 8
Collection of statistics about your cluster	Chapter 9
Error messages in the system log files	Appendix A

[Appendix A](#) lists the error messages that are generated by the Foundation Services. The appendix also provides a brief summary of the possible cause of error. The error messages are grouped by those produced during installation, those produced during runtime, and those produced when using the command-line interface. If you receive an error message, use [Appendix A](#) to find the cause and a possible resolution.

If you cannot resolve a problem that you encounter while installing or using the Foundation Services, contact your customer support center.

Using Administration Tools and Configuration Files to Troubleshoot

The Foundation Services provide a suite of tools and configuration files for maintaining and debugging your cluster environment. Because the Foundation Services run on the Solaris operating system, the Solaris tools and configuration files are also useful.

When your cluster produces an error, use the tools to determine the cause of the problem. For information about these tools, see the *Netra High Availability Suite Foundation Services 2.1 6/03 Cluster Administration Guide*.

For information about the individual tools, see their man pages in section 1M of the *Netra High Availability Suite Foundation Services 2.1 6/03 Reference Manual*. Similarly, for information about the individual configuration files, see their man pages in section 4 of the *Netra High Availability Suite Foundation Services 2.1 6/03 Reference Manual*.

Troubleshooting at Installation and Startup

For information about how to diagnose and solve problems encountered while installing and starting up the Foundation Services, see the following chapters:

- [Chapter 3](#) describes how to diagnose and solve problems during installation of the Solaris operating system and the Foundation Services.
- [Chapter 4](#) describes what to do if the master-eligible nodes do not startup correctly after a seemingly successful installation.
- [Chapter 5](#) describes what to do when you have started up the master-eligible nodes, but are unable to startup the diskless nodes or dataless nodes.

Recovering From Installation Problems

For information about how to resolve problems during installation of the Solaris operating system and the Foundation Services, see the following sections. These sections include information about how to resolve problems related to Solaris JumpStart™.

- “Recovering From General Installation Problems” on page 21
- “Recovering From `nhinstall` Problems” on page 22

Recovering From General Installation Problems

This section describes how to recover from an error that is not related to the `nhinstall` tool.

Incorrect Software Is Installed on the Cluster Nodes

If incorrect software is installed on the nodes when you use the `nhinstall` tool, perform the following procedure.

▼ **To Investigate Why Incorrect Software Is Installed on the Cluster Nodes**

1. **Search the console of the installation server for the IP address of the boot server.**

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

The IP address of the boot server in this example is the class C address 10.101.1.253.

Alternatively, use the name of the boot server to find the boot server IP address.

2. **Find the IP address of your installation server.**
3. **Compare the address of the boot server with that of the installation server.**

If the addresses are not the same, your node is being booted by the wrong machine. Perform the following steps:

 - a. **Access the console of the boot server.**
 - b. **In the `/etc/ethers` file, delete the MAC address of any nodes in your cluster.**
 - c. **In the `/etc/bootparams` file, delete the node parameters of any nodes in your cluster.**
 - d. **Restart the installation.**
4. **If you have not solved this problem, it is possibly because you have two installation servers for this cluster.**

Confirm the presence of a second installation server using the `/usr/sbin/bpgetfile` command on a machine running on the same local network as your installation server.

If the problem persists, contact your customer support center.

Recovering From `nhinstall` Problems

This section describes how to recover from error scenarios that can occur during installation using the `nhinstall` tool.

The `nhinstall` tool installs the Solaris operating system and the Foundation Services on a cluster. If the `nhinstall` tool encounters an error, it issues a message and stops. When the `nhinstall` tool encounters an error, it does not continue to search for other errors. You must fix the error and relaunch the `nhinstall` tool from the point at which it failed. If the `nhinstall` tool encounters another error, it stops again.

The `nhinstall` Tool Stops During Installation

If the `nhinstall` tool stops during the installation of a node, perform the following procedure.

▼ To Investigate Why the `nhinstall` Tool Stops During Installation

1. **Identify the problem by using the error message displayed on the installation server.**

Possible problems are:

- `local` command failure on the installation server
- `rsh` command failure due to a connection or permission problem
- `remote` command failure

2. **Correct the problem.**

If the `nhinstall` tool stops due a problem with Solaris JumpStart, see [“Solaris JumpStart Installation Fails During `nhinstall` Installation”](#) on page 23.

3. **Restart the `nhinstall` tool with the same options that you used to launch it the first time:**

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

The `nhinstall` tool resumes from the point at which it stopped.

If you have modified the `cluster_definition.conf` file to correct the error, the `nhinstall` tool displays a warning that the configuration has changed. If the change that you made to the file makes the cluster incoherent, you must reset the installation for a new installation. For information, see the *Netra High Availability Suite Foundation Services 2.1 6/03 Custom Installation Guide*.

4. **If the `nhinstall` tool stops again, repeat [Step 1](#) through [Step 3](#) until the tool completes successfully.**
5. **If you cannot resolve this problem, contact your customer support center.**

Solaris JumpStart Installation Fails During `nhinstall` Installation

If the Solaris JumpStart installation of a master-eligible node fails, perform the following procedure.

▼ To Investigate Why the JumpStart Fails During `nhinstall` Installation

1. **Confirm that the installation server and the first network interface of the master-eligible node are connected to the same switch.**

- If they are, go to [Step 2](#).
- If they are not, do the following:

a. **Configure your installation server as described in the *Netra High Availability Suite Foundation Services 2.1 6/03 Hardware Guide*.**

b. **Restart the installation.**

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

If the Solaris JumpStart software does not restart, delete the /tmp/.install_client.lock file.

2. **Confirm that a router is not connected between the installation server and the master-eligible node.**

■ If no router connects the installation server and the master-eligible node, go to [Step 3](#).

■ If a router connects the installation server and the master-eligible node, do the following:

a. **Remove the router.**

b. **Reconfigure your hardware as described in the *Netra High Availability Suite Foundation Services 2.1 6/03 Hardware Guide*.**

c. **Restart the installation:**

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

If the Solaris JumpStart software does not restart, delete the /tmp/.install_client.lock file.

3. **If you cannot resolve this problem, contact your customer support center.**

Recovering From Startup Problems on Master-Eligible Nodes

If you have installed the Foundation Services, but are unable to start up the master-eligible nodes, see the following sections:

- “A Master-Eligible Node Does Not Boot” on page 25
- “A Master Node Is Not Elected at Startup” on page 27
- “Two Master Nodes Are Elected at Startup” on page 29
- “The Vice-Master Node Remains Unsynchronized After Startup” on page 31
- “A Monitored Daemon Fails Causing a Master-Eligible Node to Reboot at Startup” on page 33
- “The Node Management Agent on a Master-Eligible Node Exits at Startup” on page 34

A Master-Eligible Node Does Not Boot

If a master-eligible node does not boot after installation, the cause could be one of the following problems:

- Incorrect hardware configuration
- Incorrect Solaris operating system configuration
- Incorrect Foundation Services configuration

If the Solaris operating system does not start on a master-eligible node, use the error messages and the Solaris documentation set to resolve the problem. If the Foundation Services do not start on a master-eligible node, perform the following procedure.

▼ To Investigate Why the Foundation Services Do Not Start on a Master-Eligible Node

1. Stop the continuous reboot cycle if such a cycle is running:

a. Access the console of the failing node.

b. Type the following command:

```
# halt  
ok>
```

Alternatively, type the following command:

```
# Control-]
telnet> send brk
Type 'go' to resume
ok>
```

The ok prompt is returned.

c. Become superuser:

```
ok> boot -s  
#
```

2. Search the error messages on the console of the failing node for an indication of the problem.

The error messages should indicate the cause of the error. For a summary of error messages and their possible causes, see [Appendix A](#).

If the error is a configuration error, the following message is displayed:

```
Error in configuration
```

The text following the message should indicate the type of configuration error. Verify that the configuration of the `nhfs.conf` file for the node is consistent with the information in the `nhfs.conf(4)` man page.

3. Confirm that the `/etc/opt/SUNWcgha/not_configured` file does not exist on the failing node.

- If the file does not exist, go to [Step 4](#).
- If the file exists, delete it and reboot the node:

```
# init 6
```

4. If the Watchdog Timer is enabled, confirm that it is configured correctly.

a. On the console of the failing node, get the ok prompt:

b. Confirm that the `nhfs.conf` file contains the parameter `WATCHDOG.NhasWatchdog=true`.

If `WATCHDOG.NhasWatchdog=false`, go to [Step 5](#).

- c. **Confirm that you have installed the correct hardware watchdog packages on the node.**

For information about the required packages and patches, see the *Netra High Availability Suite Foundation Services 2.1 6/03 README*.

- d. **Confirm that the value of the `WATCHDOG.OsTimeout` parameter in the `nhfs.conf` file is not too low.**

If you suspect that the `WATCHDOG.OsTimeout` parameter is too low, increase the value of the parameter.

- e. **Reboot the node:**

```
ok> boot
#
```

5. **If your hardware includes an OpenBoot™ PROM `diag-switch`, confirm that it is set to false:**

- a. **On the console of the failing node, get the `ok` prompt:**

- b. **Run:**

```
ok> printenv diag-switch?
```

- If the `diag-switch` is set to false, go to [Step 6](#).
- If the `diag-switch` is set to true, set it to false and reboot the node:

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

6. **If you cannot resolve this problem, contact your customer support center.**

A Master Node Is Not Elected at Startup

At startup, the first master-eligible node that you boot should become the master node. The second master-eligible node that you boot should become the vice-master node. If the first master-eligible node does not become the master node, perform the following procedure.

▼ To Investigate Why a Master Node Is Not Elected at Startup

1. **Log in to the first master-eligible node as superuser.**
2. **Confirm that the `/etc/opt/SUNWcgha/not_configured` file does not exist.**

- If the file does not exist, go to [Step 3](#).
- If the file exists, delete it and reboot the node:

```
# init 6
```

3. Confirm that the `target.conf` file has the attribute flag set to `"-"`.

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

This attribute flag indicates that a master-eligible node is qualified to become the master node. The `target.conf` file contains the node description saved by the `nhcmmd` daemon on the master node. When a master node exists, or when the cluster is running, do not edit the `target.conf` file.

- If the attribute flag is set to `"-"`, go to [Step 4](#).
- If the attribute flag is not set to `"-"`, do the following:

a. Go in to single user mode:

```
# init s
```

b. Edit the `target.conf` file to set the attribute flag to `"-"`.

The node can be set with more than one attribute flag. Make sure that the flag `"-"` is the only flag that is set.

c. Reboot the node:

```
# init 6
```

4. Confirm that the node has write access to the `cluster_nodes_table` file.

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

- If the node has write access to the file, go to [Step 5](#).
- If the node does not have write access to the file, do the following:

a. Change the access permissions as described in the `chmod(1)` man page.

b. Reboot the node:

```
# init 6
```

5. In the `cluster_nodes_table` file, confirm that the node attribute flag is set to `"-"`.

The `"-"` attribute flag indicates that the node is qualified to become the master node.

- If the attribute flag is set to `"-"`, go to [Step 6](#).
- If the attribute flag is not set to `"-"`, do the following:

a. Go in to single user mode:

```
# init s
```

b. Edit the `cluster_nodes_table` file to set the flag to `"-"`.

The attribute can be set with more than one attribute flag. Confirm that the attribute flag “-” is the only flag that is set.

c. Reboot the node:

```
# init 6
```

6. If you cannot resolve this problem, contact your customer support center.

Two Master Nodes Are Elected at Startup

At startup, when the first master-eligible node becomes the master node, the second master-eligible node should become the vice-master node. If the second master-eligible node cannot detect the master node, it will become the master node. The presence of two master nodes is an error scenario called *split brain*.

A direct link between the master-eligible nodes prevents the occurrence of split brain when the communication between the master node and vice-master node fails. For information about the direct link, see the *Netra High Availability Suite Foundation Services 2.1 6/03 Overview*.

If your cluster is configured to use a direct link, perform the procedure in [“To Investigate Split Brain on Clusters With a Direct Link”](#) on page 29. If your cluster is not configured to use a direct link, perform the procedure in [“To Investigate Split Brain on Clusters Without a Direct Link”](#) on page 30.

▼ To Investigate Split Brain on Clusters With a Direct Link

- 1. Confirm that the direct link is physically connected to the serial ports of both master-eligible nodes.**
- 2. Confirm that the `nhfs.conf` file contains the following parameters:**

```
Cluster.Direct-Link.Backend=serial
Cluster.Direct-Link.Heartbeat=20
Node.Direct-Link.serial.Device=/dev/term/b
Node.Direct-Link.serial.Speed=115200
```

The `Cluster.Direct-Link.Heartbeat` can have values up to one second.

The `Node.Direct-Link.serial.Speed` can have one of the following values: 38400, 57600, 76800, or 115200.

3. If the direct link is connected and configured correctly, and you still have a split brain error, contact your customer support center.

▼ To Investigate Split Brain on Clusters Without a Direct Link

1. Access the consoles of the master nodes.

2. Confirm that you have two master nodes.

On the console of each master-eligible node, run:

```
# nhcmmstat -c all
```

Each master node should see itself as master, and see the other master as being out of the cluster.

3. Test the communication between the master nodes.

On the console of each master-eligible node, run:

```
# nhadm check starting
```

When this command is run on a node, the command pings all of the other nodes in the cluster. If one eligible node cannot ping the other eligible node, the nodes are not communicating.

If the Carrier Grade Transport Protocol (CGTP) is installed, the `nhadm check` command pings both of the network interfaces and the CGTP interface. If CGTP is not installed, the `nhadm check` command pings one network interface only.

4. Evaluate the result obtained in [Step 3](#) by using the following table.

TABLE 4-1 Results of `nhadm check starting` Run When Two Master Nodes Are Elected at Startup

Results of <code>nhadm check</code>	Possible Cause	Action
Two network interface cards (NICs) fail, or one NIC fails and one NIC passes	Incorrect switch configuration or incorrect cabling	Reconfigure the hardware as described in the <i>Netra High Availability Suite Foundation Services 2.1 6/03 Hardware Guide</i> .
Two NICs pass but the CGTP interface fails	Incorrect Foundation Services configuration	Examine the <code>nhfs.conf</code> and <code>cluster_nodes_table</code> files.

TABLE 4-1 Results of `nhadm check` starting Run When Two Master Nodes Are Elected at Startup (Continued)

Results of <code>nhadm check</code>	Possible Cause	Action
Two NICs and the CGTP interface pass	The master-eligible nodes exist in different domains.	Confirm that the nodes have the same values for the <code>domainid</code> parameter in the <code>nhfs.conf</code> file.

5. Confirm that all of the packages and patches are installed.

a. Access the consoles of the master-eligible nodes.

b. Display the installed packages and patches:

```
# patchadd -p
# pkginfo
```

c. Compare the list of installed packages and patches with the lists defined in the *Netra High Availability Suite Foundation Services 2.1 6/03 README* and patch READMEs.

- If a required package or patch is not installed on the master-eligible node, install it and reboot both master-eligible nodes.
- If all of the required packages and patches are installed, go to [Step 6](#).

6. If you cannot resolve this problem, contact your customer support center.

The Vice-Master Node Remains Unsynchronized After Startup

After the startup of the master node and vice-master node, the data on the master node is copied to the vice-master node. In this way, the master node and vice-master node are synchronized. If the master node and vice-master node are not synchronized after startup, perform the following procedure.

▼ To Investigate Why the Vice-Master Node Remains Unsynchronized After Startup

1. Confirm that your cluster has a valid master node and vice-master node.

a. Log in to a master-eligible node as superuser.

b. Run the `nhcmmstat` command:

```
# nhcmmstat -c all
```

The `nhcmmstat` tool displays information about the roles of the peer nodes. The peer nodes should include a master node and a vice-master node. For more information about `nhcmmstat`, see the `nhcmmstat(1M)` man page.

- If your cluster has a valid master node and vice-master node, go to [Step 2](#).
- If your cluster has no master node or vice-master node, you do not have a cluster. Verify your cluster configuration by examining the `nhfs.conf` and `cluster_nodes_table` files for configuration errors.
- If your cluster has a master node but no vice-master node, reboot the master-eligible node that is not master:

```
# init 6
```

Verify that the second master-eligible node has become the vice-master node:

```
# nhcmmstat -c all
```

2. Confirm that the master node and vice-master node are unsynchronized:

```
# /usr/opt/SUNWesm/sbin/scmadm -S -M
```

If the `scmadm` tool does not reach the replicating state, the master node and vice-master node are unsynchronized. For more information, see the `nhscmadm(1M)` man page.

3. Determine whether an `nhcrfsd` daemon is running on each master-eligible node:

```
# pgrep -x nhcrfsd
```

- If a process identifier is returned, the `nhcrfsd` daemon is running. Go to [Step 5](#).
- If a process identifier is not returned, the `nhcrfsd` daemon is not running. Perform the procedure in [“To Recover From Daemon Failure” on page 52](#).

4. On the master node and vice-master node, verify that the mount point is set correctly.

The mount point is set by the `RNFS.Share` property in the `/etc/opt/SUNWcgha/nhfs.conf` file. If the mount point is set correctly, the `usr`, `root`, and `swap` parameters in the `RNFS.Share` property have the following access permissions, respectively: `ro`, `rw`, and `rw`.

5. For each node, confirm that the IP address of the `cgtp0` interface is specified in the `/etc/hosts` file.

6. If you cannot resolve this problem, contact your customer support center.

A Monitored Daemon Fails Causing a Master-Eligible Node to Reboot at Startup

When a monitored daemon fails, the Daemon Monitor triggers a recovery response. The recovery response is often to restart the failed daemon. If the daemon fails to restart correctly, the Daemon Monitor reboots the node. The failure of a monitored daemon is the most common cause of a node reboot.

If the system recovers correctly, the daemon core and error message might be the only evidence of the failure. You must take the failure seriously even though the system has recovered.

For a list of recovery responses made by the Daemon Monitor, see the `nhpmd(1M)` man page.

For information about how to recover from the failure of a monitored daemon, see [“To Recover From Daemon Failure” on page 52](#).

[Table 4–2](#) summarizes some causes of daemon failure during the startup of master-eligible nodes.

TABLE 4–2 Causes of Daemon Failure at Startup of Master-Eligible Nodes

Failed Daemon	Possible Causes at Startup
<code>nhcrfsd</code>	One of the following files on the master node contains errors: <code>/etc/vfstab</code> , <code>cluster_nodes_table</code> , or <code>nhfs.conf</code> . The local file system of the failing node is mounted or unmounted incorrectly. The network interface of the failing node is incorrectly configured.
<code>nhcmmmd</code>	One of the following files on the master node contains errors: <code>cluster_nodes_table</code> or <code>nhfs.conf</code> . The <code>cgtp0</code> interface of the failing node is incorrectly configured. The <code>cgtp0</code> interface of the failing node could not be initialized. The failing node cannot connect to the <code>nhprobed</code> daemon. The failing node cannot access the <code>/etc/services</code> file. The failing node cannot write to the <code>cluster_nodes_table</code> file when it is to be elected as master node.

TABLE 4-2 Causes of Daemon Failure at Startup of Master-Eligible Nodes (Continued)

Failed Daemon	Possible Causes at Startup
nhprobed	The failing node cannot obtain information about the network interfaces. The failing node cannot access the <code>/etc/services</code> file. The failing node cannot create the required threads, sockets, or pipe.
nhwdtd	The failing node does not have a required platform-specific plugin for the <code>nhwdtd</code> daemon. The failing node does not have a platform-specific package for hardware watchdog support. Platform-specific hardware watchdog does not work on the failing node.
in.dhcpd	A datastore location does not exist. A datastore location is not mounted on the failing node. The failing node cannot find the <code>dhcptab</code> file in the datastore.
nhnsmd	The <code>nhfs.conf</code> file on the master node contains errors.

The Node Management Agent on a Master-Eligible Node Exits at Startup

The following procedure describes what to do if the Node Management Agent (NMA) exits during the startup of the master-eligible nodes.

▼ To Investigate Why the NMA on a Master-Eligible Node Exits at Startup

1. **Confirm that the Java™ Dynamic Management Kit connector has an allocated server port number.**
 - If the server port number is already allocated, go to [Step 2](#).
 - If the server port number is not already allocated, do the following:
 - a. **In the `nma.properties` file on each peer node, allocate a port number for the Java Dynamic Management Kit connector.**
Ensure that the port number is unique.
 - b. **Restart the NMA on all nodes:**

```
# /etc/opt/SUNWcgha/init.d/nma stop
```

If the NMA fails to restart, see “NMA Not Restarted After Failure” on page 60.

2. Examine the system log files for the following messages:

- a. If the log files contain the following message, confirm that the `/etc/services` file contains an entry for the `cmm-api`.**

```
CMM statistics (JNI). Unable to access CMM statistics
(can't access cmm-api service port number).
```

- b. If the log files contain the following message, correct the `/etc/netconfig` configuration.**

```
CMM statistics (JNI). Unable to access CMM statistics
(can't access tcp netconfig).
```

- c. If the log files contain the following message, an RPC error occurred during an access to the CMM statistics.**

```
CMM statistics (JNI) Failed to get stats from CMM : [rpc return code]
Use the RPC return code to diagnose and correct the problem.
```

- d. If the log files contain the following message, a call to the CMM succeeded from an RPC point of view. However, the CMM internals were unable to return valid statistics.**

```
CMM statistics (JNI) Failed to get stats from CMM : [CMM status]
Check the status of the nhcmmmd daemon and its processes.
```

- e. If the log files contain the following message, RPC failed while attempting to access CMM statistics.**

```
CMM statistics (JNI) rpc call failed
Correct the RPC configuration.
```

- f. If the log files contain the following message, CGTP is unavailable:**

```
KSTAT (JNI). Unable to launch CGTP. CGTP statistics not available.
Confirm that the redundant network is available and that the network
configuration is correct.
```

3. Restart the NMA on all nodes:

```
# /etc/opt/SUNWcgha/init.d/nma stop
```

If the NMA fails to restart, see “NMA Not Restarted After Failure” on page 60.

4. If you cannot resolve this problem, contact your customer support center.

Recovering From Startup Problems on Diskless Nodes and Dataless Nodes

If you have started up the master-eligible nodes but are unable to startup the diskless nodes or dataless nodes, see the following sections:

- “A Diskless Node Does Not Boot at Startup” on page 37
- “A Dataless Node Does Not Boot at Startup” on page 41
- “A Monitored Daemon Fails Causing a Diskless Node or Dataless Node to Reboot at Startup” on page 42

A Diskless Node Does Not Boot at Startup

Diskless nodes boot using the Solaris Dynamic Host Configuration Protocol (DHCP) servers provided by the Reliable Boot Service. To boot diskless nodes you must have a cluster containing at least one master-eligible node running the Foundation Services. When you are booting diskless nodes, you can use the `snoop` utility to see the parameters transmitted by the DHCP server to the diskless node.

This section describes what to do when the Solaris operating system or Foundation Services do not start on a diskless node.

▼ To Investigate Why the Solaris Operating System Does Not Start on a Diskless Node

Use this procedure when the Solaris operating system does not start on a diskless node.

1. **Confirm that the spanning tree protocol is disabled.**

For Cisco 29x0 switches, do the following:

a. **Telnet to the Ethernet switch.**

b. **Type the following command:**

```
# enable
Password <user-password>
```

c. **Type the following command:**

```
# show run
```

d. **Search the output on the console for the following line:**

```
no spanning-tree vlan <vlanid>
```

If the display contains this line, the spanning tree is disabled.

- If the spanning tree is disabled, go [Step 2](#).
- If the spanning tree is not disabled, disable it.

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

2. **Confirm that the DHCP configuration is correct.**

a. **Access the consoles of the master node and vice-master node.**

b. **On each console, confirm that the `/etc/inet/dhcpsvc.conf` file exists and has the correct attributes:**

```
# nhadm check configuration
DAEMON_ENABLED=TRUE
RUN_MODE=server
RESOURCE=SUNWnhrbs
PATH=/SUNWcgha/remote/var/dhcp
CONVER=1
INTERFACE=nic0, nic1
```

The RESOURCE parameter must be set to RESOURCE=SUNWnhrbs. By default, this parameter is set to RESOURCE=SUNWfiles.

The PATH parameter must point to a directory in a replicated file system. By default, the directory is PATH=/SUNWcgha/remote/var/dhcp.

If the file does not have the correct attributes, do the following:

- Edit or create the `/etc/inet/dhcpsvc.conf` file, setting the attributes as stated previously.
- Stop and restart the DHCP daemon:

```
# /etc/rc3.d/HA.S34dhcp stop
# /etc/rc3.d/HA.S34dhcp start
```

c. **If you are installing your cluster manually, confirm that the path has DHCP container files with the following name:**

SUNWnhrbs1_10_x_1_0, SUNWnhrbs1_10_x_2_0, and SUNWnhrbs1_dhcptab, where x is the domain identity.

You do not need to perform this step if you are installing your cluster using the `nhinstall` tool.

If the DHCP container files do not have the specified name, regenerate them, taking care to use the correct values for `subnet1` and `subnet2`. For information, see “To Configure DHCP for a Diskless Node” in the *Netra High Availability Suite Foundation Services 2.1 6/03 Custom Installation Guide*.

- d. **If you are using a static address assignment, confirm that the MAC address or client ID of the diskless node is configured correctly.**

Refer to the DHCP table on the master node.

3. **On the console of the diskless node, confirm that the following OpenBoot PROM parameter is set:**

```
boot-device net:dhcp,,,,,5 net2:dhcp,,,,,5
```

4. **Confirm that the vendor type of the diskless node is recognized by the master node.**

- a. **Access the console of the master node.**

- b. **Type the following command:**

```
# snoop -v -d nic0 ether mac-adr-of-dl-node | grep -i dhcp
or
ok> dev
ok> .properties
=> property
```

The vendor type of the diskless node is returned as a string.

- c. **Search for the same string in the DHCP table on the master node.**

5. **On the console of the master node, confirm that the directory `/tftpboot` is present.**

If this directory is not present, the following error message is written to the system log files:

```
Timeout waiting for BOOTP/DHCP reply. Retrying...
TFTP Error Access violation
```

If the `/tftpboot` directory is not present on the vice-master node, the diskless node does not boot after a switchover. To set up the `/tftpboot` directory on the vice-master node, see the *Netra High Availability Suite Foundation Services 2.1 6/03 Custom Installation Guide*.

6. **On the console of the master node, confirm that the following directory contains a file for each diskless node interface:**

```
diskless_file_system/root/diskless_nodeid/etc/
```

The files could be named as follows:

```
hostname.hme0  
hostname.hme1
```

If the directory does not contain a file for an interface, the interface cannot be configured.

7. Examine the access permissions of root, swap, and usr in the `nhfs.conf` file on the master node.

If your cluster was installed manually or by the `nhinstall` tool, confirm that the following access permissions are set:

```
share -F nfs -o rw,root=diskless_node_id-nic0:diskless_node_id-nic1:  
    diskless_node_id-cgtp0 /export/root/diskless_node_id  
share -F nfs -o rw,root=diskless_node_id-nic1:diskless_node_id-nic1:  
    diskless_node_id-cgtp0 /export/swap/diskless_node_id  
share -F nfs -o ro /export/exec/Solaris_X_sparc.all/usr
```

where `X` is the version of the Solaris operating system installed on the cluster.

8. If you cannot resolve this problem, contact your customer support center.

▼ To Investigate Why the Foundation Services Do Not Start on a Diskless Node

Use this procedure when the Solaris operating system has started on the diskless node, but the Foundation Services do not start. In this error scenario, the diskless node might be in a continuous reboot cycle.

1. Stop the continuous reboot cycle if such a cycle is running:

a. Access the console of the failing node.

b. Type the following command:

```
# halt  
ok>
```

Alternatively, type the following command:

```
# Control-]
telnet> send brk
Type 'go' to resume
ok>
```

The `ok` prompt is returned.

c. Boot in single user mode:

```
ok> boot -s  
#
```

2. Search the messages displayed on the console of the failing node for an indication of the problem.

The error messages should indicate the cause of the problem. Use the error messages to identify the failing daemon or failing service. For a summary of error messages and their possible causes, see [Appendix A](#).

If the error is a configuration error, the following message is displayed:

```
Error in configuration
```

The text following the message should indicate the type of configuration error. Verify that the configuration of the `nhfs.conf` file for the node is consistent with the information in the `nhfs.conf(4)` man page.

3. **Confirm that the `/etc/opt/SUNWcgha/not_configured` file does not exist on the failing node.**
 - If the file does not exist, go to [Step 4](#).
 - If the file exists, delete it and reboot the node:

```
# init 6
```
4. **Confirm that the `cluster_node_table` file on the master node contains an entry for the failing node.**
 - If the file contains an entry for the failing node, go to [Step 5](#).
 - If the file does not contain an entry for the failing node, verify the installation and configuration.
5. **If you cannot resolve this problem, contact your customer support center.**

A Diskless Node Does Not Boot After Failover

If a failover occurs during the boot or reboot of a diskless node, the DHCP files can be corrupted. If this problem occurs, see [“A Diskless Node Does Not Reboot After Failover”](#) on page 48.

A Dataless Node Does Not Boot at Startup

Dataless nodes boot from a local disk and run customer applications locally. Dataless nodes access the Foundation Services through the cluster network and send data to the master node.

This section describes what to do if the Solaris operating system or the Foundation Services do not start on a dataless node. Only use this section when you have a running cluster that contains a master node and a vice-master node.

If the Solaris operating system does not start on a dataless node, use the error messages and the Solaris documentation set to resolve the problem. If the Foundation Services do not start on a dataless node, perform the following procedure.

▼ To Investigate Why the Foundation Services Do Not Start on a Dataless Node

1. **Stop the continuous reboot cycle if such a cycle is running.**

For information, see [Step 1 of “To Investigate Why the Foundation Services Do Not Start on a Diskless Node”](#) on page 40.

2. **Search the messages on the console of the failing node for an indication of the problem.**

For information, see [Step 2 of “To Investigate Why the Foundation Services Do Not Start on a Diskless Node”](#) on page 40.

3. **Confirm that the `/etc/opt/SUNWcgha/not_configured` file does not exist on the failing node.**

For information, see [Step 3 of “To Investigate Why the Foundation Services Do Not Start on a Diskless Node”](#) on page 40.

4. **Confirm that the `cluster_node_table` file on the master node contains an entry for the failing node.**

For information, see [Step 4 of “To Investigate Why the Foundation Services Do Not Start on a Diskless Node”](#) on page 40.

5. **If you cannot resolve this problem, contact your customer support center.**

A Monitored Daemon Fails Causing a Diskless Node or Dataless Node to Reboot at Startup

When a monitored daemon fails, the Daemon Monitor triggers a recovery response. The recovery response is often to restart the failed daemon. If the daemon fails to restart correctly, the Daemon Monitor reboots the node. The failure of a monitored daemon is the most common cause of a node reboot.

If the system recovers correctly, the daemon core and error message might be the only evidence of the failure. You must take the failure seriously even though the system has recovered.

For a list of recovery responses made by the Daemon Monitor, see the `nhpmd(1M)` man page.

For information about how to recover from the failure of a monitored daemon, see [“To Recover From Daemon Failure” on page 52](#).

[Table 5–1](#) summarizes some causes of daemon failure during the startup of diskless nodes and dataless nodes.

TABLE 5–1 Causes of Daemon Failure on Diskless Nodes and Dataless Nodes at Startup

Failed Daemon	Possible Causes at Startup
<code>nhcmmd</code>	<p>One of the following files on the master node contains errors: <code>cluster_nodes_table</code> or <code>nhfs.conf</code>.</p> <p>The <code>cgtp0</code> interface of the failing node is configured incorrectly.</p> <p>The <code>cgtp0</code> interface of the failing node could not be initialized.</p> <p>The failing node cannot connect to the <code>nhprobed</code> daemon.</p> <p>The failing node cannot access the <code>/etc/services</code> file.</p> <p>The failing node exceeded the time-out value.</p>
<code>nhprobed</code>	<p>The failing node cannot obtain information about the network interfaces.</p> <p>The failing node cannot access the <code>/etc/services</code> file.</p> <p>The failing node cannot create the required threads, sockets, or pipe.</p>
<code>nhwdtd</code>	<p>The failing node does not have a required platform-specific plugin for the <code>nhwdtd</code> daemon.</p> <p>The failing node does not have a platform-specific package for hardware watchdog support.</p> <p>Platform-specific hardware watchdog support does not work on the failing node.</p>

Troubleshooting at Runtime

For information about how to diagnose and solve problems encountered after a successful startup, see the following chapters:

- [Chapter 6](#) describes how to recover from problems associated with failover and switchover.
- [Chapter 7](#) describes how to solve problems that cause a node to reboot during runtime.
- [Chapter 8](#) describes what to do if you cannot add a node to your cluster.
- [Chapter 9](#) describes error scenarios that can occur when you are using the NMA.

Recovering From Failover and Switchover Problems

For information about how to recover from problems associated with failover and switchover, see the following sections:

- [“Two Master Nodes Are Elected at Runtime” on page 47](#)
- [“Replication Does Not Resume After Failover or Switchover” on page 49](#)

Two Master Nodes Are Elected at Runtime

During runtime, one master-eligible node should be the master node, and the other master-eligible node should be the vice-master node. When both master-eligible nodes act as master nodes, you have an error scenario called split brain. For information about split brain and the use of a direct link, see [“Two Master Nodes Are Elected at Startup” on page 29](#).

If a split brain error occurs during runtime on a cluster with a direct link, perform the procedure in [“To Investigate Split Brain on Clusters With a Direct Link” on page 29](#). If a split brain error occurs during runtime on a cluster without a direct link, perform the procedure in [“To Investigate Split Brain During Runtime on Clusters Without a Direct Link” on page 47](#).

▼ To Investigate Split Brain During Runtime on Clusters Without a Direct Link

1. **Access the consoles of the master nodes.**
2. **Confirm that you have two master nodes.**

On the console of each master-eligible node, run:

```
# nhcmmstat -c all
```

Each master node should see itself as master, and see the other master as being out of the cluster.

3. Test the communication between the master nodes.

On the console of each master-eligible node, run:

```
# nhadm check starting
```

When this command is run on a node, it pings all of the other nodes in the cluster. If one master node cannot ping the other master node, the nodes are not communicating.

- If the nodes are able to communicate with each other, go to [Step 4](#).
- If the nodes are not able to communicate with each other, examine the network interface values of the nodes.

For information, see “Examining the Cluster Networking Configuration” in the *Netra High Availability Suite Foundation Services 2.1 6/03 Cluster Administration Guide*.

When the problem is resolved, Reliable NFS should automatically detect the split-brain situation. Reliable NFS reboots the master-eligible nodes, so that there is a master node and a vice-master node.

4. Determine whether the spanning tree protocol is disabled, as described in [Step 1](#) of “[To Investigate Why the Solaris Operating System Does Not Start on a Diskless Node](#)” on page 37.

5. If you cannot resolve this problem, contact your customer support center.

A Diskless Node Does Not Reboot After Failover

If a failover occurs during the boot or reboot of a diskless node, the DHCP files can be corrupted. If this problem occurs, perform the following procedure.

▼ To Reboot a Diskless Node After Failover

1. Confirm that the cluster has recovered from the failover.

For information, see “Reacting to a Failover” in the *Netra High Availability Suite Foundation Services 2.1 6/03 Cluster Administration Guide*.

2. Reconfigure the boot policy for the diskless node.

For information, see “Configuring DHCP for a Diskless Node” in the *Netra High Availability Suite Foundation Services 2.1 6/03 Custom Installation Guide*.

3. Reload the DHCP table on the master node and the vice-master node:

```
# kill -HUP in.dhcpd
```

Replication Does Not Resume After Failover or Switchover

If replication does not resume after failover or switchover, examine the replication between the master-eligible nodes, as described in [“The Vice-Master Node Remains Unsynchronized After Startup”](#) on page 31.

Recovering From Node Reboot at Runtime

For information about the causes of node reboot at runtime, see [“A Monitored Daemon Fails Causing a Node to Reboot at Runtime”](#) on page 51.

A Monitored Daemon Fails Causing a Node to Reboot at Runtime

When a monitored daemon fails, the Daemon Monitor triggers a recovery response. The recovery response is often to restart the failed daemon. If the daemon fails to restart correctly, the Daemon Monitor reboots the node. The failure of a monitored daemon is the most common cause of a node reboot.

If the system recovers correctly, the daemon core and error message might be the only evidence of the failure. You must take the failure seriously even though the system has recovered.

For a list of recovery responses made by the Daemon Monitor, see the `nhpmd(1M)` man page. For a summary of the causes of daemon failure during startup, see [“A Monitored Daemon Fails Causing a Master-Eligible Node to Reboot at Startup”](#) on page 33 and [“A Monitored Daemon Fails Causing a Diskless Node or Dataless Node to Reboot at Startup”](#) on page 42.

[Table 7-1](#) and [Table 7-2](#) summarize the events that can cause a monitored daemon to fail at runtime. To recover from daemon failure, perform the procedure in [“To Recover From Daemon Failure”](#) on page 52.

TABLE 7-1 Possible Causes of Daemon Failure at Runtime

Failed Daemon	Possible Cause at Runtime
nhcmmmd	The nhcmmmd daemon was killed. The failing node does not see its presence in the <code>cluster_nodes_table</code> .
nhprobed	The nhprobed daemon was killed.
nhwdtd	The operating system has hung. The system is overloaded.

TABLE 7-2 Causes of Daemon Failure on Master-Eligible Nodes During Failover or Switchover

Failed Daemon	Possible Cause During Failover or Switchover
nhcrfsd	The nhcrfsd daemon was killed during the failover or switchover.
nhcmmmd	The node cannot connect to the nhprobed daemon.
nhprobed	The node cannot create the required threads, sockets, or pipe.

▼ To Recover From Daemon Failure

1. Examine the core file produced by the failed daemon.

The core file is located in the `/var/tmp/SUNWcgha/core` directory, and has the format:

core.node_name.executable_file_name.process_ID.time

For more information about core dumps, see the `coreadm(1M)` man page.

2. Examine the system log files for an error message produced by the failed daemon.

For example, the following error message is produced by the failure of a daemon launched by the `rpc` nametag:

```
[ID 615790 local0.notice] "rpc" Failed to stay up.
```

For information about which nametag launches which daemon, see the `nhpmd(1M)` man page.

3. Identify the cause of the daemon failure.

Use the information obtained in [Step 1](#), [Step 2](#), [Table 7-1](#), and [Table 7-2](#).

4. Fix the underlying problem, if necessary.

5. Confirm that the recovery procedure has been carried out by searching the system log files for `local0` information.

- If your system log file is not configured for local0 information, reconfigure it. For information, see the *Netra High Availability Suite Foundation Services 2.1 6/03 Cluster Administration Guide*.
- If local0 information is logged to a file, search the file for the string “nhpmd”. Lines containing the string “nhpmd” describe the recovery response performed by the Daemon Monitor.

Cannot Add Nodes to a Running Cluster

When you have a running cluster of two master-eligible nodes, you can add diskless nodes and dataless nodes. For information about supported cluster configurations, see the *Netra High Availability Suite Foundation Services 2.1 6/03 Hardware Guide*. For information about how to add a node to a running cluster, see the *Netra High Availability Suite Foundation Services 2.1 6/03 Cluster Administration Guide*.

If you cannot add a node to a cluster, read the following section.

Cannot Add a Node to a Running Cluster by Using the `nhinstall` Tool

For information about how to add diskless nodes to a cluster by using the `nhinstall` tool, see the *Netra High Availability Suite Foundation Services 2.1 6/03 Cluster Administration Guide*. If the node or nodes are not added successfully, perform the following procedure.

Note – You cannot add dataless nodes to a cluster by using the `nhinstall` tool.

▼ To Investigate Why You Cannot Add a Diskless Node to a Running Cluster by Using the `nhinstall` Tool

1. **Confirm that you are able to add nodes to your cluster without generating an unsupported cluster configuration.**

For information about supported cluster configurations, see the *Netra High Availability Suite Foundation Services 2.1 6/03 Hardware Guide*.

a. Log in to a peer node.

b. Establish your current cluster configuration:

```
# nhcmmstat -c all
```

The `nhcmmstat` tool displays information about the nodes in the cluster. For information, see the `nhcmmstat(1M)` man page.

c. Confirm that you will not generate an unsupported cluster configuration by adding your new nodes.

- If your new cluster configuration is supported, go to [Step 2](#).
- If your new cluster configuration is not supported, you cannot add the new nodes.

2. Verify that your cluster has a Solaris environment for diskless nodes.

On a master-eligible node, look for a directory called `/export/os`, where `os` is `Solaris_8` or `Solaris_9`.

- If the environment exists, go to [Step 3](#).
- If the environment does not exist, do the following:

a. Create the Solaris environment.

For information, see the *Netra High Availability Suite Foundation Services 2.1 6/03 Custom Installation Guide*.

b. Add the new diskless node.

3. Verify that the nodes that you want to install are defined in the `cluster_definition.conf` file and are used in the `nhinstall add` command.

For example, the `cluster_definition.conf` file for a cluster of two master-eligible nodes would contain the following lines:

```
NODE=10 08:00:20:f9:b3:6a
NODE=20 08:00:20:f9:aa:66
```

To add two diskless nodes, you edit the `cluster_definition.conf` file to add the two new nodes. For example, to add node 30 and node 40, edit the file as follows:

```
NODE=10 08:00:20:f9:b3:6a
NODE=20 08:00:20:f9:aa:66
NODE=30 -                               - node30
NODE=40 -                               - node40
```

When you run the `nhinstall add` command, identify *both* of the nodes:

```
# /opt/SUNWcgha/sbin/nhinstall -r /home/nhasconf add 30 40
```

where `/home/nhasconf` is the directory containing the `nhinstall` configuration files.

The installation fails if you identify one node only, for example:

```
# /opt/SUNWcgha/sbin/nhinstall -r /home/nhasconf add 30
```

If the installation fails, rerun the `nhinstall add` command and identify all of the new nodes that are defined in the `cluster_definition.conf` file.

4. Determine whether you have changed any options in the `cluster_definition.conf` file since you installed the master-eligible nodes.

The `cluster_definition.conf` file must contain the same options as it had at the initial installation of the master-eligible nodes.

5. Verify that the `addon.conf` file, if present, is configured correctly.

For information about the `addon.conf` file, see the `addon.conf(4)` man page.

6. If you still cannot add diskless nodes to the cluster, reinstall your cluster. Include the diskless nodes in your new cluster definition.

Cannot Collect Statistics by Using the Node Management Agent

For information about how to recover from errors that occur when you are using the Node Management Agent (NMA), see the following sections:

- “An External Client Cannot Communicate With the Node Management Agent” on page 59
- “NMA Not Restarted After Failure” on page 60
- “NMA Not Sending SNMP Traps to a Given Target” on page 61
- “The `switchOver` Method Does Not Finish Executing” on page 61
- “Cascading Fails” on page 61

An External Client Cannot Communicate With the Node Management Agent

If an external client cannot communicate with the NMA on a peer node, perform the following procedure.

▼ To Investigate Why an External Client Cannot Communicate With an NMA on a Peer Node

1. **Confirm that an RMI registry has not been instantiated on a designated NMA TCP port.**

When a Java Dynamic Management Kit agent enables an RMI connector, the agent instantiates an RMI registry on the port it is using. If an RMI registry uses the same TCP port as the NMA, communication with the NMA is impossible. The RMI registry that initially allocated the port has control of the port.

2. **If the NMA is not accessible through RMI, do the following.**

- a. Uncomment the `java.rmi.server` property in the `nma.properties` file.
- b. Set the value of the `java.rmi.server` property to the IP address or host name of the node.
- c. Restart the NMA on all nodes:

```
# /etc/opt/SUNWcgha/init.d/nma stop
```

If the NMA fails to restart, see “NMA Not Restarted After Failure” on page 60.

3. If you cannot resolve this problem, contact your customer support center.

NMA Not Restarted After Failure

The NMA process is monitored by the Daemon Monitor. If the NMA on a peer node fails, the Daemon Monitor restarts it. If the NMA fails to restart, perform the following procedure.

▼ To Investigate Why the NMA Is Not Restarted

1. Find the current number of times that the Daemon Monitor has attempted to restart the NMA.

This parameter is called the *reset retry count*. For information, see “Daemon Monitor Statistics” in the *Netra High Availability Suite Foundation Services 2.1 6/03 NMA Programming Guide*.

2. Confirm that the reset retry count does not exceed 10.

The Daemon Monitor attempts to restart the NMA up to 10 times.

- If the reset retry count has not been exceeded, go to [Step 3](#).
- If the reset retry count has been exceeded, try one or both of the following:

- a. Reset the reset retry counter by using the `resetRetryCount` method.

For information, see “Manipulating Daemon Monitor Retry Settings” in the *Netra High Availability Suite Foundation Services 2.1 6/03 NMA Programming Guide*.

- b. Analyze the traces to diagnose the cause of the problem.

3. If you cannot resolve this problem, contact your customer support center.

NMA Not Sending SNMP Traps to a Given Target

If the NMA does not send SNMP traps to a target, confirm that the following files are configured:

- `installDir/etc/opt/SUNWcgha/nma.targets.txt`
- `installDir/etc/opt/SUNWcgha/nma.params.txt`
- `installDir/etc/opt/SUNWcgha/nma.notifs.txt`

where `installDir` is the root installation directory. Replace `installDir` with the root installation directory if the root installation directory is not `/`.

For information about these files, see their man pages in the *Netra High Availability Suite Foundation Services 2.1 6/03 Reference Manual*. For information about sending SNMP traps, see “Developing a Remote SNMP Manager” in the *Netra High Availability Suite Foundation Services 2.1 6/03 NMA Programming Guide*.

The `switchOver` Method Does Not Finish Executing

The `switchOver` method is run on the NMA on the master node to initiate a switchover. When you use the `switchOver` method, do not use the floating external address to access the master node. When the floating external address is used, the `switchOver` method breaks the connection. Instead, use a fixed IP address on the master node.

Cascading Fails

Cascading is the transfer of statistics from the NMA on each of the peer nodes into the namespace of the master node. Cascading provides the NMA of the master node with a view of the entire cluster. If the NMA on a peer node cannot be detected by the NMA on the master node, perform the following procedure.

▼ To Examine Why Cascading Fails

1. Confirm that cascading is enabled.

a. Open the `/etc/opt/SUNWcgha/nma.properties` file.

b. Confirm that the following line is commented out:

```
com.sun.nhas.ma.cascading.enabled=false
```

■ If the line is commented out, go to [Step 2](#).

■ If the line is not commented out, comment it out and restart the NMA:

```
# /etc/opt/SUNWcgha/init.d/nma stop
```

If the NMA fails to restart, see [“NMA Not Restarted After Failure”](#) on page 60.

2. Confirm that the maximum number of HTTP clients has not been exceeded.

a. Connect to the NMA of the node that is not cascading correctly, using an Internet browser.

For information, see *“Accessing the NMA”* in the *Netra High Availability Suite Foundation Services 2.1 6/03 NMA Programming Guide*.

b. Find the number of connected clients.

The number of connected clients is specified in the

```
com.sun.nhas.ma.connectors.http.client
```

 property.

c. Find the maximum number of connected clients.

■ Open the `/etc/opt/SUNWcgha/nma.properties` file.

■ Find the value of the `connectors.http.client` parameter.

d. If the number of connected clients exceeds the maximum, increase the maximum.

■ Change the `connectors.http.client` parameter in the `/etc/opt/SUNWcgha/nma.properties` file.

■ Restart the NMA:

```
# /etc/opt/SUNWcgha/init.d/nma stop
```

If the NMA fails to restart, see [“NMA Not Restarted After Failure”](#) on page 60.

3. If you cannot resolve this problem, contact your customer support center.

Error Messages

This appendix lists the error messages that are generated by the Foundation Services and provides a brief summary of the possible cause of each error. The error messages are grouped by those produced during installation, those produced during runtime, and those produced when using the command-line interface. Each list appears in alphabetical order of error message.

For information about the error messages that are produced by the Foundation Services, see the following sections:

- “Introduction to Error Messages” on page 63
- “Error Messages Written During Installation” on page 64
- “Error Messages Written During Runtime” on page 65
- “Error Messages Written by Command-Line Tools” on page 70

Introduction to Error Messages

All of the Foundation Services log error and information messages to the system log files. These messages can be processed by client programs or by the Node Management Agent. For information about how to configure system log files, see the *Netra High Availability Suite Foundation Services 2.1 6/03 Cluster Administration Guide*.

This appendix lists some of the error messages that are produced by the error scenarios described in this book. This appendix does not contain an exhaustive list of error messages produced by the Foundation Services.

All error messages have the following format:

```
<date><nodeid><PID><stack><failed service or daemon><message>
```

The following output shows some sample error messages:

```
1:Nov  8 10:05:20 b14-netra-4 CMM[547]: [ID 191875 local0.error]
    S-CMM [Membership] (no_master_role) Cluster becomes stale
3:Oct 24 15:47:20 b14-netra-4 statd[3162]: [ID 514559 daemon.error]
    svc_tp_create: Could not register prog 100024 vers 1 on udp
3:Oct 24 15:47:20 b14-netra-4 nhcrfsd[591]: [ID 191875 local0.error]
    S-CORE nhcrfsd' failed to stay up
3:Oct 24 15:47:25 b14-netra-4 nfs: [ID 609386 kern.warning] WARNING:
    lockd: cannot contact statd (error 4), continuing
3:Oct 24 15:53:43 b14-netra-4 CMM[546]: [ID 191875 local0.error]
    S-CMM [PROBE] Node 173 becomes DOWN
```

Error Messages Written During Installation

Error Messages Written During Manual Installation

```
[nhcmm] /etc/opt/SUNWcgha/not_configured present => no action
taken
```

The installation or configuration process has failed.

```
[nhwdt] /etc/opt/SUNWcgha/not_configured present => no action
taken
```

The installation or configuration process has failed.

Error Messages Written During Installation Using the `nhinstall` Tool

The `nhinstall` tool is a script that runs a series of commands. If the `nhinstall` tool encounters an error during the installation, it issues a message to the console and stops. The error message indicates which command has failed. If the failing command is a Solaris command, the `nhinstall` tool indicates which Solaris command has failed. The `nhinstall` tool does not produce its own error messages.

When the `nhinstall` tool encounters an error, it does not continue to search for other errors. You must fix the error, then relaunch the `nhinstall` tool from the point at which it failed. If the `nhinstall` tool encounters another error, it stops again.

Error Messages Written During Runtime

Error Messages Written by the Cluster Membership Manager

- [%s] port undefined in /etc/services
Confirm that the `nhfs.conf` and `cluster_nodes_table` files are correctly configured. See the man pages in the *Netra High Availability Suite Foundation Services 2.1 6/03 Reference Manual*.
- [%s/%s] port undefined in /etc/services
Confirm that the `nhfs.conf` and `cluster_nodes_table` files are correctly configured. See the man pages in the *Netra High Availability Suite Foundation Services 2.1 6/03 Reference Manual*.
- Another CMM is running => exit
A second instance of the `nhcmm` daemon attempted to start on the node.
- [Config] %s does not exist, could not recreate %s
A cluster configuration file or its backup is missing. The file cannot be created. Confirm that the directory in which the file should exist has the correct access permissions.
- [Config] Could not initialize Cluster nodes table (local)
Confirm that the `cluster_nodes_table` file is correctly configured. See the `cluster_nodes_table(4)` man page.
- [Config] Could not initialize Cluster nodes table (NFS)
Confirm that the `cluster_nodes_table` file is correctly configured. See the `cluster_nodes_table(4)` man page.
- [Config] Could not initialize minimal config
Confirm that the `target.conf` file is correctly configured. See the `target.conf(4)` man page.
- [Config] Could not load backup configuration file
A cluster configuration file or its backup is missing. The file cannot be created. Confirm that the directory in which the file should exist has the correct access permissions.
- [Config] Could not read Cluster nodes table
Investigate whether the `cluster_nodes_table` file is missing or inaccessible.
- [Config] Could not read minimal configuration file
Investigate whether the `target.conf` file is missing or inaccessible.

- [Config] Invalid domain id in file %s
Confirm that the `nhfs.conf` and `cluster_nodes_table` files are correctly configured. See the man pages in the *Netra High Availability Suite Foundation Services 2.1 6/03 Reference Manual*.
- [Config] Invalid node id in file %s
Confirm that the `nhfs.conf` and `cluster_nodes_table` files are correctly configured. See the man pages in the *Netra High Availability Suite Foundation Services 2.1 6/03 Reference Manual*.
- [Config] Invalid node name in file %s
Confirm that the `cluster_nodes_table` files is correctly configured. See the `cluster_nodes_table(4)` man page.
- [Config] Minimal configuration files %s and %s cannot be accessed
Investigate whether the `target.conf` file is missing or inaccessible.
- [Candidates] No master-eligible node
See [Chapter 4](#).
- [Config] Nodes number exceed cluster capacities
For information about the supported cluster configuration, see the *Netra High Availability Suite Foundation Services 2.1 6/03 Hardware Guide*
- CURRENT NODE HAS LEFT CLUSTER
The current node does not belong to the master node domain. Confirm that the following are true:
- The node `domainid` is correct in the `nhfs.conf` file and the `cluster_nodes_table` file.
 - The node name is correct in the `cluster_nodes_table` file.
 - The node `domainid` is the same in the `nhfs.conf` file and the `cluster_nodes_table` file.
 - The master node `domainid` is correct.
- [Election] Unable to extract the best choice
No node is eligible to become the master node. Verify that the vice-master node is qualified, and that it is synchronized with the master node. See [“The Vice-Master Node Remains Unsynchronized After Startup” on page 31](#).
- [Membership] Could not access Cluster nodes table
Investigate whether the `cluster_nodes_table` file is missing or inaccessible.
- [PROBE] a probe is already executed
Confirm whether the following are true:
- The node is present in the `cluster_nodes_table` file.
 - The node name is unique in the `cluster_nodes_table` file.
 - The `nodeid` is unique in the `cluster_nodes_table` file.

Error Messages Written by Reliable NFS

Already mounted. In order to ensure data integrity, please proceed to a 'full sync' using 'nhcrfsadm -f %s' command

This error occurs if a partition is already mounted at boot time.

Could not disable previous SNDR configuration

Reliable NFS could not disable the SNDR boot time configuration. This can happen if the replication configuration is broken. Flush the replication configuration manually by performing the following steps:

1. Boot the master-eligible nodes in single-user mode:

```
ok> boot -s
```

2. Reset the replication configuration on both nodes:

```
# /usr/opt/SUNWscm/sbin/dscfg -i
```

3. Re-create an empty replicated configuration file by typing Y at this prompt on both nodes:

```
# (Type Y for YES) Y
```

```
# /usr/opt/SUNWscm/sbin/dscfg -i -p /etc/opt/SUNWesm/pconfig
```

4. Reboot the nodes.

5. If the problem persists:

- a. Boot both master-eligible nodes in single user mode.

- b. On each master-eligible node, edit the `/etc/opt/SUNWcgha/target.conf` file by setting the attributes field to "-" and the role field to "-".

For information about the `target.conf` file, see the `target.conf(4)target.conf(4)` man page.

- c. Repeat Steps 2 to 4 on each master-eligible node.

Could not export some directories

Reliable NFS could not share some directories. Verify that the directories listed to be shared exist and that `/usr/bin/share` exists.

Could not get port number for server *<port number>*

Reliable NFS could not bind to the specified service. Verify that the service is not defined in the `/etc/services` file. If it is not defined, add an entry to the `/etc/services` file to define the service.

Could not put SNDR into logging mode

You cannot stop SNDR. Examine the disk configuration.

Could not reverse SNDR configuration

You cannot reverse SNDR during a switchover. This role reversal is handled by switching the primary and secondary SNDR roles. Examine the disk configuration.

Could not set master dynamic address(es)
 Reliable NFS could not set the master node floating address triplet. Verify that the interfaces exist.

Could not start <command name>
 Reliable NFS could not execute the specified command. Verify that the command is available on the cluster and that its execution rights are correct.

Could not start SNDR
 The SNDR service failed to start. Verify that the node has a valid disk partition configuration. For information about disk partitions, see the *Netra High Availability Suite Foundation Services 2.1 6/03 Custom Installation Guide*.

Could not stop <command name>
 Reliable NFS could not execute the specified command. Verify that the command is available on the cluster, and that its execution rights are correct.

Could not unexport some directories
 Reliable NFS could not share some directories. Verify that the directories listed to be shared exist, and that /usr/bin/share exists.

Could not unset master dynamic address(es)
 Reliable NFS could not unset the master node floating address triplet. The specified interfaces might be unknown or unplumbed.

Emergency reboot of the node
 Reliable NFS rebooted the node because it did not restart correctly. This can occur if the nhcrfsd daemon dies during a switchover or a failover.

Error in configuration
 The Reliable NFS configuration is incorrect. The text following the message should indicate the type of configuration error. Verify that the configuration of the nhfs.conf file for the failing node is consistent with the nhfs.conf(4) man page.

Illegal startup case: we are 'master' but were 'vice-master unsynchronized'. Please restart the node.
 The vice-master node was rebooted and became the master node. This scenario is not allowed. Nodes must be restarted with the same role as they had before shutdown.

Mount of local filesystems failed
 Reliable NFS could not mount or unmount local filesystems and abort. Verify that mount points and file systems are coherent. For the device listed in the error, check:

- The associated mount point in the /etc/vfstab file
- The access permission of this mount point

No canonical name found for address <IP address>
 No canonical name was found to correspond to the specified address. Canonical names are required for every address. Specify the canonical name for this address in /etc/hosts.

Node's CMM mastership and RNFS one is not coherent
 Reliable NFS believes that the current node is the master, but the nhcmm daemon does not consider the current node to be the master.

Number of SNDR slices is greater than configuration file one
SNDR slices are configured but not managed through Reliable NFS.

Unable to read kstat data

A partition managed by Reliable NFS disappeared while the cluster was running. This might happen if you change the SNDR configuration while the cluster is running. This scenario is not allowed. Reboot the node.

Unmount of local filesystems failed

See "Mount of local filesystems failed".

Vice master has <number> slices, we have <number>: refusing vice master to follow

The master disk and vice-master disks do not have the same disk partition configuration. This is not allowed. Stop the vice-master and change its disk partition configuration to be the same as that on the master. See "Modifying and Adding Disk Partitions" in the *Netra High Availability Suite Foundation Services 2.1 6/03 Cluster Administration Guide*.

Vice master has a wrong configuration: refusing vice master to follow

The master's view of the disk configuration on the vice-master disk is not current. Check that the `nhfs.conf` file is consistent between the two nodes.

Wrong slice configured in SNDR

This problem occurs if SNDR slices are configured but not managed through Reliable NFS. Do not use SNDR on behalf of Reliable NFS.

Error Messages Written by the Reliable Boot Service

`cmm_connect() failed (#)`

Examine log files for messages from `nhpmd` saying that the `nhcmmnd` daemon was stopped. If necessary, reboot the node to restart the `nhcmmnd` daemon.

Error Messages Written by the Watchdog Timer

CPCI: cannot find 'watchdog-level1' node in PICL tree

Verify the patch level of the `SUNWpiclu` package against the required patches listed in the release notes of the hardware platform.

CPCI: configure: cannot connect to PICL daemon

The `picl` daemon is not running. Reboot the node to restart this daemon.

LOM: cannot stat /dev/lom

Verify that the LOM driver packages, `SUNWlommu` and `SUNWlomr`, are installed. If these packages are not installed, install them.

nhwdtd could not read config file, exiting

The nhwdtd daemon cannot find the `nhfs.conf` file, or the contents of this file are invalid. Compare the contents of the file with the requirements described in the `nhfs.conf(4)` man page.

Error Messages Written by the Node Management Agent

CMM statistics (JNI) Failed to get stats from CMM: [CMM status]

A call to the CMM succeeded from an RPC point of view. However, the CMM internals were unable to return valid statistics. Check the status of the `nhcmm` daemon and its processes.

CMM statistics (JNI) Failed to get stats from CMM: [rpc return code]

An RPC error occurred during an access to the CMM statistics. Use the RPC return code to diagnose and correct the problem.

CMM statistics (JNI). Unable to access CMM statistics (can't access `cmm-api` service port number)

The CMM is incorrectly configured. Confirm that `/etc/services` contains an entry for `cmm-api`.

CMM statistics (JNI). Unable to access CMM statistics (can't access `tcp netconfig`).

The `netconfig` database is incorrectly configured for TCP. Correct the `/etc/netconfig` configuration.

CMM statistics (JNI) rpc call failed

RPC failed while attempting to access Cluster Membership Manager statistics. Correct the RPC configuration.

KSTAT (JNI). Unable to launch CGTP

CGTP statistics are not available. Confirm that the redundant network is available and that the network configuration is correct.

Error Messages Written by Command-Line Tools

This section contains the error messages written by the Foundation Services command-line tools. For information about these tools, see their man pages in section 1M of the *Netra High Availability Suite Foundation Services 2.1 6/03 Reference Manual*.

Error Messages Written by the nhadm Command

`<file>` does not exist

Ensure that the path to the file is correct.

`<interface>` (`<IP address>`) has no associated host name in `/etc/hosts`

The `<IP address>` corresponding to the `<interface>` is not defined in `/etc/hosts`.

`<interface>`: invalid IP address (`<current>` instead of `<expected>`)

The IP address configured for the specified interface is not correct.

`<IP address>` defined twice in `/etc/hosts`

Confirm that only one hostname is defined for the specified IP address.

`<node1>` and `<node2>` have the same `node_id`

Both nodes have the same `nodeid` in the `cluster_nodes_table`. Ensure that all nodes have a unique `nodeid`.

`<parameter>`: `<current>` instead of `<expected>`

Set the correct value for the specified parameter in the `/etc/inet/dhcpsvc.conf` file.

`<parameter>` not defined

The specified parameter must be defined in the `/etc/inet/dhcpsvc.conf` file.

`<service>` service: `<current>` instead of `<expected>`

Correct the protocol for the specified service in `/etc/services`.

`<service>` service must be defined in `/etc/services`

Add an entry for the specified service in `/etc/services`. Daemons using an unspecified service might not work correctly.

`<slice>` is not replicated

The slice specified in the CRFS configuration file is incorrectly replicated.

`<slice>` not defined in `/etc/vfstab`

All slices defined in the CRFS configuration file must have an entry in the `/etc/vfstab` file.

`<slice>`: `<bitmap slice>` does not look like a raw device

A bitmap slice must be a raw device.

`<slice>`: `<bitmap slice>` does not exist

A bitmap slice must exist for each slice that is to be replicated. Create the bitmap slice with the correct size, or verify its name.

can not determine root file system

No root partition was found when executing the `mount` command.

can not read `<file>`: permission denied

Ensure that root user can access the specified file.

CGTP interface is not present

The Foundation Services are probably not being started.

Destination `<CGTP ip>`, gateway `<interface IP>`: no route defined
 To be routed. CGTP protocol requires a route to be defined for each interface.

`/etc/hostname.<interface>` does not exist
 This file is required and should contain the name of the node for the specified interface.

`"forcedirectio"` mount option set for root file system `<slice>`
 This option must not be set for the specified slice.

`"forcedirectio"` mount option set for slice `<slice>`
 This option must not be set for the specified slice.

host name not defined for `<interface>`
 The `/etc/hostname` file does not exist or does not contain a host name.

ill-formatted line (`<line number>`)
 The line contained in the file being processing is incorrect. Check its syntax.

incorrect domain id for node `<node name>`
 The domainid for the specified node is not equal to the value indicated in the CMM configuration file. The domainid is specified by the `CMM_DOMAIN_ID` variable.

invalid host name for `<interface>` (`<IP address>`): (`<current>` instead of `<expected>`)
 The host name specified for `<interface>` does not match the following:

- The name found in the `/etc/hostname.<interface>` file for a physical interface
- The name found in the `cluster_nodes_table` file for the CGTP interface

Invalid kernel mode: `<bits>`-bit instead of 64-bit mode
 Install the correct version of the Solaris operating system.

Invalid node id for node `<node>`: `<node id>` reserved for master logical address
 The specified nodeid is reserved for the master node floating address. Use another value for nodeid.

local node (id=`<node_id>`) not defined in the `cluster_nodes_table`
 The local node has no entry in the `cluster_nodes_table` file. Add an entry for this node.

`"logging"` mount option not set for root file system `<slice>`
 This option is required for the specified slice.

`"logging"` mount option not set for slice `<slice>`
 This option is required for the specified slice.

`"mount at boot"` for slice `<slice>` not set to `"no"`
 This option prevents the slice from being automatically mounted at boot. An incorrect option can cause this slice to be incorrectly managed by Reliable NFS.

Netra HA Suite services will not start
 See [Chapter 4](#) or [Chapter 5](#).

No partition for the mounting point *<name>*
 No partition is defined for the specified mounting point.

node *<node name>* *<parameter>*: invalid parameter
 The specified parameter is not known. Check the syntax of the specified file.

node *<node name>*: invalid role (*<value>*)
 The specified role is not correct.

package *<name>* not installed
 Install the specified package.

package *<name>*: version *<current>* (at least *<minimum>* required)
 Install the correct version of the specified package.

package *<name>*: version *<current>* instead of *<expected>*
 Install the correct version of the specified package.

Partition size too small for mounting point *<name>*: *<real>* sectors
 instead of expected
 Change the partition size.

Patch *<patch>* not installed
 Install the specified patch.

release *<current>* instead of *<expected>*
 Install the correct version of the Solaris operating system.

remote slice *<remote>* not equal to local slice *<local>*
 For each replicated slice specified in the CRFS configuration file, the remote slice
 name and the local slice name must be the same.

root file system *<slice>* not defined in */etc/vfstab*
 The slice corresponding to the root file system must be defined in the
/etc/vfstab file.

PATH not defined
 PATH is not defined in the */etc/inet/dhcpsvc.conf* file.

PATH does not reference a valid file (*<file>*)
 Ensure that the file referenced by PATH in the */etc/inet/dhcpsvc.conf* file
 exists, or set PATH with the correct path name.

Slice *<slice>*: replication bitmap slice *<bitmap>* too small: (*<x>*)
 bytes instead of *<y>*)
 The bitmap slice is too small to support replication of the specified slice. Increase
 the size of the bitmap slice.

symbol *<symbol>* not defined in */etc/dhcp/inittab*
 The specified symbol is required to allow information to be retrieved from the
 DHCP requests on a diskless node.

/usr/kernel/drv/rdc.conf must contain: *rdc_bitmap_mode=1*
 Add the setting of *rdc_bitmap_mode*. An incorrect value can cause Reliable NFS
 to work incorrectly.

Error Messages Written by the nhcrfsadm Command

This order is not allowed at this time

A synchronization is in progress, or no vice-master node exists.

Unable to contact services on host

The service is not available, or the nhcrfsd daemon is not be running on this node.

Index

A

adding nodes to a cluster, 55-57

B

boot

continuous reboot cycle, 26, 40, 42

dataless node does not boot at startup, 41-42

diskless node does not boot after failover, 48

diskless node does not boot at startup, 37

master-eligible node does not boot at startup, 25

C

cannot resolve a problem, 15

cascading, fails, 61

Cisco switches, confirming that the spanning tree protocol is disabled, 37

cluster

adding nodes, 55-57

cannot add nodes by using `nhinstall`, 55

confirming the Solaris environment for diskless nodes, 56

supported configurations, 55

Cluster Membership Manager, error

messages, 65-66

command-line tools, error messages, 70-74

commands, for debugging, 17

communication

testing communication between master-eligible nodes, 30, 48

configuration files, for debugging, 17

continuous reboot cycle, 26, 40, 42

customer support center, when to contact, 15

D

daemon

causes of failure at runtime, 51

causes of failure at startup of diskless and dataless nodes, 43

causes of failure at startup of master-eligible node, 33

recovering from failure, 52

Daemon Monitor

cannot restart the NMA, 60

finding and resetting the reset retry count, 60

dataless node

does not boot at startup, 41-42

Foundation Services do not start, 42

reboots because a monitored daemon fails at runtime, 51-53

reboots because a monitored daemon fails at startup, 42-43

Solaris operating system does not start, 42

DHCP

configuration, 38

container files, 38

- diskless node
 - debug using `snoop`, 37
 - does not boot after failover, 48
 - does not boot at startup, 37
 - Foundation Services do not start, 40
 - reboots because a monitored daemon fails at runtime, 51-53
 - reboots because a monitored daemon fails at startup, 42-43
 - Solaris environment for, 56
 - Solaris operating system does not start, 37
- documentation, related documents, 8

E

- election
 - no master node at startup, 27
 - two master nodes at runtime, 47-48
 - two master nodes at startup, 29
- endless reboot, recovery from, 67
- error messages
 - Cluster Membership Manager, 65-66
 - command-line tools, 70-74
 - format, 63
 - installation, 64
 - introduction to, 63-64
 - manual installation, 64
 - `nhadm`, 71-73
 - `nhcrfsadm`, 74
 - `nhinstall`, 64
 - NMA, 70
 - Reliable Boot Service, 69
 - Reliable NFS, 67-69
 - runtime, 65-70
 - Watchdog Timer, 69-70

F

- failover
 - diskless node boot fails, 48
 - two master nodes, 47-48
- FAQs, 13
- floating external address, being used by the NMA, 61
- Foundation Services
 - do not start on a dataless node, 42

- Foundation Services (Continued)
 - do not start on a diskless node, 40
 - do not start on a master-eligible node, 26

H

- HTTP clients, finding the maximum number, 62

I

- installation
 - error messages, 64
 - incorrect software installed, 21
 - `nhinstall` problems, 22
 - problems, 21
 - Solaris JumpStart fails while using `nhinstall`, 23
 - stops while using `nhinstall`, 22-23

J

- Java Dynamic Management Kit connector, allocating a server port number, 34

M

- manual installation, error messages, 64
- master-eligible node
 - does not boot at startup, 25
 - Foundation Services do not start, 26
 - reboots because a monitored daemon fails at runtime, 51-53
 - reboots because a monitored daemon fails at startup, 33-34
 - Solaris operating system does not start, 25
 - startup problems, 25-35
 - testing communication between master-eligible nodes, 30, 48
- master-ineligible node
 - reboots because a monitored daemon fails at runtime, 51-53
 - reboots because a monitored daemon fails at startup, 42-43

- master node
 - cannot synchronize at startup, 31-32
 - no master node at startup, 27
 - two master nodes at runtime, 47-48
 - two master nodes at startup, 29
- monitored daemon
 - causes of failure at runtime, 51
 - causes of failure at startup of diskless and dataless nodes, 43
 - causes of failure at startup of master-eligible node, 33
 - fails at runtime, 51-53
 - fails at startup of diskless and dataless nodes, 42-43
 - fails at startup of master-eligible node, 33-34
 - recovering from failure, 52
- mount point, verifying, 32

N

- nhadm, error messages, 71-73
- nhadm check starting
 - testing communication between master-eligible nodes, 30, 48
- nhcmmstat
 - finding the cluster configuration, 55
 - finding the role of a node, 30, 47
- nhcrfsadm, error messages, 74
- nhcrfsadm, resynchronizing disks, 31
- nhcrfsd, verifying that the nhcrfsd daemon is running, 32
- nhinstall
 - cannot add nodes to a cluster, 55
 - error messages, 64
 - problems, 22
 - Solaris JumpStart fails, 23
 - stops during installation, 22-23
- nhinstall add, adding nodes to a cluster, 56
- NMA

- allocating a server port number for the Java Dynamic Management Kit connector, 34
- cannot communicate with external client, 59
- cannot initiate a switchover, 61
- cascading fails, 61
- does not restart after failure, 60
- error messages, 70

NMA (Continued)

- exists at startup of master-eligible nodes, 34-35
- finding the maximum number of HTTP clients, 62
- making accessible through RMI, 59
- not sending SNMP traps, 61
- restarting after failure, 34
- runtime problems, 59
- using the floating external address, 61
- node
 - adding to a cluster, 55-57
 - cannot add to a cluster by using nhinstall, 55
 - reboot at runtime, 51-53
 - reboot at startup of diskless and dataless nodes, 42-43
 - reboot at startup of master-eligible node, 33-34
 - Solaris environment for diskless nodes, 56

O

- OpenBoot PROM, configuration of, 27

Q

- questions, commonly asked, 13

R

- reboot
 - at runtime, 51-53
 - at startup of diskless and dataless nodes, 42-43
 - at startup of master-eligible node, 33-34
 - continuous reboot cycle, 26, 40, 42
 - diskless node does not reboot after failover, 48
- related documents, 8
- Reliable Boot Service, error messages, 69
- Reliable NFS
 - error messages, 67-69
 - verifying that the nhcrfsd daemon is running, 32

- reset retry count, finding and resetting, 60
- RNFS.Share* property, entry in *nhfs.conf*, 32
- role
 - finding the role of a node, 30, 47, 55
- runtime
 - causes of node reboot, 51-53
 - error messages, 65-70
 - monitored daemon fails, 51-53
 - NMA problems, 59
 - problems with failover and switchover, 47-49
 - two master nodes, 47-48

S

- scmadm*, verifying synchronization, 31
- SNMP
 - configuration files, 61
 - traps not being sent by the NMA, 61
- snoop*, 37
- software, incorrect software installed, 21
- Solaris environment, for diskless nodes, 56
- Solaris JumpStart, fails while using *nhinstall*, 23
- Solaris operating system
 - does not start on a dataless node, 42
 - does not start on a diskless node, 37
 - does not start on a master-eligible node, 25
- spanning tree protocol, confirming that it is disabled, 37
- split brain
 - at runtime, 47
 - at startup, 29
- startup
 - cannot synchronize master node and vice-master node, 31-32
 - dataless node does not boot, 41-42
 - diskless node does not boot, 37
 - Foundation Services do not start on a dataless node, 42
 - Foundation Services do not start on a diskless node, 40
 - Foundation Services do not start on a master-eligible node, 26
 - master-eligible node does not boot, 25
 - monitored daemon fails on diskless and dataless nodes, 42-43

- startup (Continued)
 - monitored daemon fails on master-eligible node, 33-34
 - NMA exists at startup of master-eligible nodes, 34-35
 - no master node, 27
 - of master-eligible nodes, 25-35
 - problems with installation, 21-24
 - Solaris operating system does not start on a dataless node, 42
 - Solaris operating system does not start on a diskless node, 37
 - Solaris operating system does not start on a master-eligible node, 25
 - two master nodes, 29
- switchover
 - cannot be initiated by the NMA, 61
 - two master nodes, 47-48
- switchOver* method, does not finish executing, 61
- synchronization
 - cannot synchronize at startup, 31-32
 - verifying mount point, 32

T

- tools, for debugging, 17

V

- vice-master node
 - cannot synchronize at startup, 31-32
 - two master nodes at runtime, 47-48
 - two master nodes at startup, 29

W

- Watchdog Timer
 - configuration, 26
 - error messages, 69-70