



Netra™ High Availability Suite 3.0 1/08 Foundation Services Troubleshooting Guide

Sun Microsystems, Inc.
www.sun.com

Part No. 819-5248-13
March 2008, Revision A

Submit comments about this document at: <http://www.sun.com/hwdocs/feedback>

Copyright 2008 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.sun.com/patents>, and one or more additional patents or pending patent applications in the U.S. and in other countries.

This document and the product to which it pertains are distributed under licenses restricting their use, copying, distribution, and decompilation. No part of the product or of this 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, Java, docs.sun.com, 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.

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.

[IF ENERGY STAR INFORMATION IS REQUIRED FOR YOUR PRODUCT, DO THE FOLLOWING: DELETE THIS TEXT. DOWNLOAD THE ENERGY STAR GRAPHIC (ENERGYSTAR.EPS) FROM DOCS MANAGER TO YOUR /ART DIRECTORY. USE THE <GRAPHIC> ELEMENT TEXT.]

U.S. Government Rights—Commercial use. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

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 2008 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, Californie 95054, Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. a les droits de propriété intellectuelle relatants à la technologie qui est décrit dans ce document. En particulier, et sans la limitation, ces droits de propriété intellectuelle peuvent inclure un ou plus des brevets américains énumérés à <http://www.sun.com/patents> et un ou les brevets plus supplémentaires ou les applications de brevet en attente dans les Etats-Unis et dans les autres pays.

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 des systèmes 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, Java, docs.sun.com, et Solaris sont des marques de fabrique ou des marques déposées 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.

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 licenciées de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

LA DOCUMENTATION EST FOURNIE "EN L'ÉTAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISÉE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.



Adobe PostScript

Contents

Preface ix

1. Using the Troubleshooting Guide 1

When to Use This Book 1

How to Use This Book 2

Using Administration Tools and Configuration Files to Troubleshoot 3

2. Recovering From Installation Problems 5

Recovering From General Installation Problems 5

Incorrect Software Is Installed on the Cluster Nodes 5

- ▼ To Investigate Why Incorrect Software Is Installed on the Cluster Nodes 6

Recovering From `nhinstall` Problems 7

The `nhinstall` Tool Stops During Installation 7

- ▼ To Investigate Why the `nhinstall` Tool Stops During Installation 7

Solaris JumpStart Installation Fails During `nhinstall` Installation 8

- ▼ To Investigate Why the JumpStart Fails During `nhinstall` Installation 8

3. Recovering From Startup Problems on Master-Eligible Nodes 11

A Master-Eligible Node Does Not Boot 11

▼ To Investigate Why the Netra HA Suite Does Not Start on a Master-Eligible Node	12
A Master Node Is Not Elected at Startup	13
▼ To Investigate Why a Master Node Is Not Elected at Startup	13
Two Master Nodes Are Elected at Startup	15
▼ To Investigate Split Brain on Clusters With a Direct Link	15
▼ To Investigate Split Brain on Clusters Without a Direct Link	16
The Vice-Master Node Remains Unsynchronized After Startup	17
▼ To Investigate Why the Vice-Master Node Remains Unsynchronized After Startup	17
A Monitored Daemon Fails Causing a Master-Eligible Node to Reboot at Startup	19
The Node Management Agent on a Master-Eligible Node Exits at Startup	20
▼ To Investigate Why the NMA on a Master-Eligible Node Exits at Startup	21
4. Recovering From Startup Problems on Diskless Nodes and Dataless Nodes	25
A Diskless Node Does Not Boot at Startup	25
▼ To Investigate Why the Solaris Operating System Does Not Start on a Diskless Node	26
▼ To Investigate Why the Netra HA Suite Does Not Start on a Diskless Node	29
A Diskless Node Does Not Boot After Failover	31
A Dataless Node Does Not Boot at Startup	31
▼ To Investigate Why the Netra HA Suite Does Not Start on a Dataless Node	31
A Monitored Daemon Fails Causing a Diskless Node or Dataless Node to Reboot at Startup	32
5. Recovering From Failover and Switchover Problems	35
Two Master Nodes Are Elected at Run Time	35
▼ To Investigate Split Brain During Run Time on Clusters Without a Direct Link	35

A Diskless Node Does Not Reboot After Failover	36
▼ To Reboot a Diskless Node After Failover	37
Replication Does Not Resume After Failover or Switchover	37
6. Recovering From Node Reboot at Run Time	39
A Monitored Daemon Fails Causing a Node to Reboot at Run Time	39
▼ To Recover From Daemon Failure	40
7. Cannot Add Nodes to a Running Cluster	43
Cannot Add a Node to a Running Cluster by Using the <code>nhinstall</code> Tool	43
▼ To Investigate Why You Cannot Add a Diskless Node to a Running Cluster by Using the <code>nhinstall</code> Tool	43
8. Cannot Collect Statistics by Using the Node Management Agent	47
An External Client Cannot Communicate With the Node Management Agent	47
▼ To Investigate Why an External Client Cannot Communicate With an NMA on a Peer Node	48
NMA Not Restarted After Failure	48
▼ To Investigate Why the NMA Is Not Restarted	48
NMA Not Sending SNMP Traps to a Given Target	49
The <code>switchOver</code> Method Does Not Finish Executing	49
Cascading Fails	50
▼ To Examine Why Cascading Fails	50
A. Error Messages	53
Introduction to Error Messages	53
Error Messages Written During Installation	54
Error Messages Written During Manual Installation	54
Error Messages Written During Installation Using the <code>nhinstall</code> Tool	54
Error Messages Written During Run Time	56
Error Messages Written by the Cluster Membership Manager	56

Error Messages Written by Reliable NFS	60
Error Messages Written by the Reliable Boot Service	64
Error Messages Written by the Node Management Agent	64
Error Messages Written by Command-Line Tools	65
Error Messages Written by the <code>nhadm</code> Command	66
Error Messages Written by the <code>nhcrfsadm</code> Command	71
Index	73

Tables

TABLE 1-1	Common Questions	1
TABLE 1-2	Where to Find Troubleshooting Information	2
TABLE 3-1	Causes of Daemon Failure at Startup of Master-Eligible Nodes	20
TABLE 4-1	Causes of Daemon Failure on Diskless Nodes and Dataless Nodes at Startup	33
TABLE 6-1	Possible Causes of Daemon Failure at Run Time	40
TABLE 6-2	Causes of Daemon Failure on Master-Eligible Nodes During Failover or Switchover	40

Preface

This book describes how to solve problems when installing, starting up, and maintaining the Netra™ High Availability (HA) Suite Foundation Services 3.0 product.

Who Should Use This Book

This book is for system administrators who are installing or maintaining the Netra HA Suite. This book assists administrators with the following troubleshooting tasks:

- Diagnosing and solving problems encountered while installing, configuring, or maintaining the Netra HA Suite
 - Understanding error messages and resolving errors
-

How This Book Is Organized

This book contains the following chapters and appendixes:

[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 and information about tools and configuration files you can use to troubleshoot.

[Chapter 2](#) describes how to diagnose and solve problems during installation of the Solaris™ Operating System and the Netra HA Suite.

[Chapter 3](#) describes what to do if the master-eligible nodes do not startup correctly after a seemingly successful installation.

[Chapter 4](#) describes what to do when you have started up the master-eligible nodes but are unable to startup the diskless nodes or dataless nodes.

[Chapter 5](#) describes how to recover from problems associated with failover and switchover.

[Chapter 6](#) describes how to solve problems that cause a node to reboot during run time.

[Chapter 7](#) describes what to do if you cannot add a node to your cluster.

[Chapter 8](#) 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 Netra HA Suite and describes how to resolve the errors.

Shell Prompts

Shell	Prompt
C shell	<i>machine-name%</i>
C shell superuser	<i>machine-name#</i>
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#

Typographic Conventions

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

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

Related Documentation

The following table lists the documentation for this product. The online documentation is available at:

<http://docs.sun.com/app/docs/prod/netra.ha30>

Application	Title	Part Number
Late-breaking news	<i>Netra High Availability Suite 3.0 1/08 Release Notes</i>	819-5249-14
Introduction to concepts	<i>Netra High Availability Suite 3.0 1/08 Foundation Services Overview</i>	819-5240-13
Basic setup, supported hardware, and configurations	<i>Netra High Availability Suite 3.0 1/08 Foundation Services Getting Started Guide</i>	819-5241-13
Automated installation methods	<i>Netra High Availability Suite 3.0 1/08 Foundation Services Installation Guide</i>	819-5242-13
Detailed installation methods	<i>Netra High Availability Suite 3.0 1/08 Foundation Services Manual Installation Guide for the Solaris OS</i>	819-5237-13

Application	Title	Part Number
Cluster administration	<i>Netra High Availability Suite 3.0 1/08 Foundation Services Cluster Administration Guide</i>	819-5235-13
Using the Cluster Membership Manager	<i>Netra High Availability Suite 3.0 1/08 Foundation Services CMM Programming Guide</i>	819-5236-13
Using the SAF CMM API	<i>Netra High Availability Suite 3.0 1/08 Foundation Services SA Forum Programming Guide</i>	819-5246-13
Using the Node Management Agent	<i>Netra High Availability Suite 3.0 1/08 Foundation Services NMA Programming Guide</i>	819-5239-13
Configuring outside the cluster using CGTP	<i>Netra High Availability Suite 3.0 1/08 Foundation Services Standalone CGTP Guide</i>	819-5247-13
Man pages for Foundation Services features and APIs using the Solaris OS	<i>Netra High Availability Suite 3.0 1/08 Foundation Services Solaris Reference Manual</i>	819-5244-13
Man pages for Foundation Services features and APIs using Linux	<i>Netra High Availability Suite 3.0 1/08 Foundation Services Linux Reference Manual</i>	819-5245-12
Definitions and acronyms	<i>Netra High Availability Suite 3.0 1/08 Foundation Services Glossary</i>	819-5238-13
Common problems	<i>Netra High Availability Suite 3.0 1/08 Foundation Services Troubleshooting Guide</i>	819-5248-13

Documentation, Support, and Training

The Sun web site provides information about the following additional resources:

- Documentation (<http://www.sun.com/documentation>)
- Support (<http://www.sun.com/support>)
- Training (<http://www.sun.com/training>)

Third-Party Web Sites

Sun is not responsible for the availability of third-party web sites mentioned in this document. Sun does not endorse and is not responsible or liable for any content, advertising, products, or other materials that are available on or through such sites or resources. Sun will not be responsible or liable for any actual or alleged damage or loss caused by or in connection with the use of or reliance on any such content, goods, or services that are available on or through such sites or resources.

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can submit your comments by going to:

<http://www.sun.com/hwdocs/feedback>

Please include the title and part number of your document with your feedback:

Netra™ High Availability Suite 3.0 1/08 Foundation Services Troubleshooting Guide, part number 819-5248-13

Using the Troubleshooting Guide

- For information about when and how to use this book, and for pointers to other books in the Netra™ High Availability (HA) Suite Foundation Services documentation set, see the following sections:
- “When to Use This Book” on page 1
 - “How to Use This Book” on page 2
 - “Using Administration Tools and Configuration Files to Troubleshoot” on page 3

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 Netra HA Suite. 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 Netra HA Suite. These subjects are covered in other books in the documentation set.

The following table lists some common questions and answers about the Netra HA Suite.

TABLE 1-1 Common Questions

Question	For Information, see
What is the Netra HA Suite?	<i>Netra High Availability Suite 3.0 1/08 Foundation Services Overview</i>
What's new in this release?	<i>Netra High Availability Suite 3.0 1/08 Release Notes</i>
What tools and configuration files are available?	<i>Netra High Availability Suite 3.0 1/08 Foundation Services Reference Manual</i> for the man pages that can be accessed by using the <code>man1</code> command

TABLE 1-1 Common Questions (*Continued*)

Question	For Information, see
How can the Netra HA Suite be evaluated easily?	<i>Netra High Availability Suite 3.0 1/08 Foundation Services Getting Started Guide</i>
How is the hardware configured?	<i>Netra High Availability Suite 3.0 1/08 Foundation Services Getting Started Guide</i>
Which installation method should you use?	<i>Netra High Availability Suite 3.0 1/08 Foundation Services Overview</i>
How is the Netra HA Suite installed?	<i>Netra High Availability Suite Foundation Services 3.0 Automated Installation Guide and Netra High Availability Suite 3.0 1/08 Foundation Services Manual Installation Guide for the Solaris OS</i>
How is a cluster maintained, and how is hardware added or replaced?	<i>Netra High Availability Suite 3.0 1/08 Foundation Services Cluster Administration Guide</i>
How are applications that use the Netra HA Suite developed?	<i>Netra High Availability Suite 3.0 1/08 Foundation Services CMM Programming Guide, the Netra High Availability Suite 3.0 1/08 Foundation Services SA Forum Programming Guide, or the Netra High Availability Suite 3.0 1/08 Foundation Services NMA Programming Guide</i>
What does a term or a name mean?	<i>Netra High Availability Suite 3.0 1/08 Foundation Services 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	Chapter 3 through Chapter 5
Installation of the Solaris Operating System and the Netra HA Suite	Chapter 2
Startup of master-eligible nodes	Chapter 3
Startup of diskless nodes or dataless nodes	Chapter 4

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

Problem Area	For Information, see
Run time	Chapter 6 through Chapter 8
Failover and switchover during run time	Chapter 5
Node reboot during run time	Chapter 6
Addition of nodes to your cluster during run time	Chapter 7
Collection of statistics about your cluster	Chapter 8
Error messages in the system log files	Appendix A

[Appendix A](#) lists the error messages that are generated by the Netra HA Suite. The appendix also provides a brief summary of the possible cause of error. The error messages are grouped as follows: those produced during installation, those produced during run time, 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 Netra HA Suite, contact your customer support center.

Using Administration Tools and Configuration Files to Troubleshoot

The Netra HA Suite provides a suite of tools and configuration files for maintaining and debugging your cluster environment. Because the Netra HA Suite runs 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 3.0 1/08 Foundation Services Cluster Administration Guide*.

For information about the individual tools, see their man pages in section 1M of the *Netra High Availability Suite 3.0 1/08 Foundation Services Reference Manual*. Similarly, for information about the individual configuration files, see their man pages in section 4 of the *Netra High Availability Suite 3.0 1/08 Foundation Services Reference Manual*.

Recovering From Installation Problems

For information about how to resolve problems during installation of the Solaris Operating System and the Netra HA Suite, see the following sections. These sections include information about how to resolve problems related to Solaris JumpStart™.

- [“Recovering From General Installation Problems” on page 5](#)
- [“Recovering From `nhinstall` Problems” on page 7](#)

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. For example, for a Solaris installation:**

```
Beginning system identification...
Searching for configuration file(s)...
Using sysid configuration file
10.101.1.253:/export/nhjumpstart/101/10/sysidcfg
Search complete.
```

For a Linux installation:

```
IP-Config: Got DHCP answer from 10.101.1.253, my address is \
10.101.1.10
IP-Config: Complete:
    device=eth0, addr=10.101.1.10, mask=255.255.255.0, \
    gw=10.101.1.253, host=cp3020-1, domain=, nis-domain=(none), \
    bootserver=10.101.1.253, rootserver=10.101.1.253, \
    rootpath=/dist/mvista/CGE_4.0/target_tdp<7>eth0: no IPv6 \
    routers present
```

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. **If the boot server is running the Solaris OS, access the console of the wrong boot server as follows:**
 - i. **Delete the MAC address of any nodes in your cluster in the `/etc/ethers` file.**
 - ii. **Delete the node parameters of any nodes in your cluster in the `/etc/bootparams` file.**
 - iii. **Delete the node parameters of any nodes in your cluster in the DHCP configuration files under `/var/dhcp`**
 - iv. **Restart the installation.**

OR

- 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 8](#).

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

On the Solaris OS:

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

On Linux:

```
# /opt/sun/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 3.0 1/08 Foundation Services Manual Installation Guide for the Solaris OS*.

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 3.0 1/08 Foundation Services Getting Started 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.lck 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 3.0 1/08 Foundation Services Getting Started 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.lck 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 Netra HA Suite, but are unable to start up the master-eligible nodes, see the following sections:

- [“A Master-Eligible Node Does Not Boot” on page 11](#)
- [“A Master Node Is Not Elected at Startup” on page 13](#)
- [“Two Master Nodes Are Elected at Startup” on page 15](#)
- [“The Vice-Master Node Remains Unsynchronized After Startup” on page 17](#)
- [“A Monitored Daemon Fails Causing a Master-Eligible Node to Reboot at Startup” on page 19](#)
- [“The Node Management Agent on a Master-Eligible Node Exits at Startup” on page 20](#)

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 Netra HA Suite 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 Netra HA Suite does not start on a master-eligible node, perform the following procedure.

▼ To Investigate Why the Netra HA Suite Does Not Start on a Master-Eligible Node

1. Stop the continuous reboot cycle if such a cycle is running. For example on a SPARC node with an Open Boot PROM (OBP):

- 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 `thenhfs.conf4` man page.

3. Confirm that the `/etc/opt/SUNWcgaha/not_configured` file on Solaris (respectively `/etc/opt/sun/nhas/not_configured` on Linux) does not exist on the failing node.

- If the file does not exist, go to the next step.
- If the file exists, delete it and reboot the node.

4. 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 5](#).
- If the diag-switch is set to true, set it to false and reboot the node:

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

5. 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 on the Solaris OS (respectively `/etc/opt/sun/nhas/not_configured` on Linux) does not exist.
 - If the file does not exist, go to [Step 3](#).
 - If the file exists, delete it and reboot the node.

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

For more information, see the `target.conf4` 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 `nhcmmnd` 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.
 - 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.

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

For more information, see the `cluster_nodes_table4` 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 `chmod1` man page.
 - b. Reboot the node.

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.

- 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.

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 3.0 1/08 Foundation Services 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 15](#). 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 16](#).

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

The *Cluster.Direct-Link.Heartbeat* is expressed in seconds.

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.

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 3.0 1/08 Foundation Services Getting Started Guide</i> .
Two NICs pass but the CGTP interface fails	Incorrect Netra HA Suite configuration	Examine the <code>nhfs.conf</code> and <code>cluster_nodes_table</code> files.
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:

On the Solaris OS:

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

On Linux:

```
# rpm -qa
```

- c. Compare the list of installed packages and patches with the lists defined in the Netra High Availability Suite 3.0 1/08 Foundation Services 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 `nhcmmstat1M` 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.

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:

On Linux, run the following command:

```
# /sbin/drbdadm state all
```

If at least one partition is not in the state Primary/Secondary, the master node and vice-master node are unsynchronized. For more information, see the `drbdadm(8)` man page.

For versions of the Solaris OS earlier than version 10:

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

For versions of the Solaris 10 OS and later:

```
# /usr/sbin/dsstat 1
```

If the `scmadm` or `dsstat` tool does not reach the replicating state, the master node and vice-master node are unsynchronized. For more information, see the `nhscmadm1M` 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 40](#).
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 on the Solaris OS and on `/etc/opt/SUNWcgha/nhfs.conf` Linux. 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 dump 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 `nhpmd1M` man page.

For information about how to recover from the failure of a monitored daemon, see [“To Recover From Daemon Failure” on page 40](#).

TABLE 3-1 summarizes some causes of daemon failure during the startup of master-eligible nodes.

TABLE 3-1 Causes of Daemon Failure at Startup of Master-Eligible Nodes

Failed Daemon	Possible Causes at Startup
nhcrfsd	One of the following files on the master node contains errors: /etc/vfstab on Solaris, /etc/vfstab on Linux, cluster_nodes_table, or nhfs.conf. The local file system of the failing node is mounted or unmounted incorrectly. The network interface of the failing node is incorrectly configured.
nhcmmd	One of the following files on the master node contains errors: cluster_nodes_table or nhfs.conf. The cgtp0 interface of the failing node is incorrectly configured. The cgtp0 interface of the failing node could not be initialized. The failing node cannot connect to the nhprobed daemon. The failing node cannot access the /etc/services file. The failing node cannot write to the cluster_nodes_table file when it is to be elected as master node.
nhprobed	The failing node cannot obtain information about the network interfaces. The failing node cannot access the /etc/services file. The failing node cannot create the required threads, sockets, or pipe.
in.dhcpd	A datastore location does not exist. A datastore location is not mounted on the failing node. The failing node cannot find the dhcptab file in the datastore.
nhnsmd	The nhfs.conf 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 port numbers were correctly configured.

- Determine whether the `nma.properties` file content is consistent with the information provided in the `nma.properties(4)` man page.
- Ensure that none of the ports described in the `nma.properties` file are already in use.

For example, if the `nma.properties` file contains the following line:

```
com.sun.nhas.ma.adaptors.snmp.port=8085
```

you can check that the port is not already in use by typing the following:

```
# /bin/netstat -a | grep 8085
```

If this command returns something, it means that this port is already in use. In this case, change the value of this port in the `nma.properties` file to a value that is not yet in use.

2. Once the port numbers have been properly configured, you can start the NMA on this node by doing one of the following:

On the Solaris 9 OS:

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

On the Solaris 10 OS:

```
# /usr/sbin/svccadm restart svc:/system/cgha/nma
```

3. 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).
```

At this point, the `/etc/services` should contain the following content:

udp6	tpi_clts	v	inet6	udp	/dev/udp6	-
tcp6	tpi_cots_ord	v	inet6	tcp	/dev/tcp6	-
udp	tpi_clts	v	inet	udp	/dev/udp	-
tcp	tpi_cots_ord	v	inet	tcp	/dev/tcp	-

rawip	tpi_raw	-	inet	-	/dev/rawip	-
ticlts	tpi_clts	v	loopback	-	/dev/ticlts	
straddr.so						
ticotsord	tpi_cots_ord	v	loopback	-	/dev/ticotsord	
straddr.so						
ticots	tpi_cots	v	loopback	-	/dev/ticots	
straddr.so						

- b. 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.

- c. If the log file contains a CMM statistics (JNI) failure, like the following message:**

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

or:

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

or:

```
CMM statistics (JNI) rpc call failed
```

This means there is a problem in the connection between the NMA and the CMM. Check the status of the nhcmmd daemon and its processes.

- d. 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 by running the following command:

```
# nhadm check
```

4. Restart the NMA on all nodes as follows:

On Solaris 9 OS:

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

The NMA will then be relaunched by the PMD.

On the Solaris 10 OS:

```
# /usr/sbin/svcadm restart svc:/system/cgha/nma
```

5. If you have problems relaunching the NMA, 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 25](#)
- [“A Dataless Node Does Not Boot at Startup” on page 31](#)
- [“A Monitored Daemon Fails Causing a Diskless Node or Dataless Node to Reboot at Startup” on page 32](#)

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 Netra HA Suite. 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 Netra HA Suite does 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 3.0 1/08 Foundation Services Getting Started 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:

```
# cat /etc/inet/dhcpsvc.conf
DAEMON_ENABLED=TRUE
RUN_MODE=server
RESOURCE=SUNWnhrbs
PATH=/SUNWcgga/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=/SUNWcgga/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.

If your cluster is running the Solaris 9 OS, issue the following commands:

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

If your cluster is running the Solaris 10 OS, issue the following commands:

```
# svcadm restart svc:/network/dhcp-server:default
```

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 about configuring DHCP for a diskless node, see the *Netra High Availability Suite 3.0 1/08 Foundation Services Manual Installation Guide for the Solaris OS*.

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. For SPARC, on the console of the diskless node, confirm that the following OpenBoot PROM parameter is set:

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

Note – For information about performing this task on x64 platforms, refer to the products hardware documentation.

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 3.0 1/08 Foundation Services Manual Installation Guide for the Solaris OS*.

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.hme0hostname.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_idshare -F nfs -o
rw,root=diskless_node_id-nic1:diskless_node_id-nic1:
    diskless_node_id-cgtp0 /export/swap/diskless_node_idshare -F nfs -o ro
/export/exec/Solaris_X_sparc.all/usr
```

where X is the version of the Solaris OS installed on the cluster.

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

▼ To Investigate Why the Netra HA Suite Does Not Start on a Diskless Node

Use this procedure when the Solaris Operating System has started on the diskless node, but the Netra HA Suite 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  
#
```

Note – For information about performing this task on x64 platforms, refer to the products hardware documentation.

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.conf4` 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.

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 36.](#)

A Dataless Node Does Not Boot at Startup

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

This section describes what to do if the Solaris Operating System or the Netra HA Suite 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 Netra HA Suite does not start on a dataless node, perform the following procedure.

▼ To Investigate Why the Netra HA Suite Does 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 Netra HA Suite Does Not Start on a Diskless Node” on page 29.](#)

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 Netra HA Suite Does Not Start on a Diskless Node” on page 29.](#)

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

For information, see [Step 3 of “To Investigate Why the Netra HA Suite Does Not Start on a Diskless Node” on page 29.](#)

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 Netra HA Suite Does Not Start on a Diskless Node” on page 29.](#)

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 dump 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 `nhpmd1M` man page.

For information about how to recover from the failure of a monitored daemon, see [“To Recover From Daemon Failure” on page 40.](#)

TABLE 4-1 summarizes some causes of daemon failure during the startup of diskless nodes and dataless nodes.

TABLE 4-1 Causes of Daemon Failure on Diskless Nodes and Dataless Nodes at Startup

Failed Daemon	Possible Causes at Startup
nhcmmnd	One of the following files on the master node contains errors: cluster_nodes_table or nhfs.conf.
	The cgtpp0 interface of the failing node is configured incorrectly.
	The cgtpp0 interface of the failing node could not be initialized.
	The failing node cannot connect to the nhprobed daemon.
	The failing node cannot access the /etc/services file.
	The failing node exceeded the time-out value.
nhprobed	The failing node cannot obtain information about the network interfaces.
	The failing node cannot access the /etc/services file.
	The failing node cannot create the required threads, sockets, or pipe.

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 Run Time” on page 35](#)
- [“Replication Does Not Resume After Failover or Switchover” on page 37](#)

Two Master Nodes Are Elected at Run Time

During run time, 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 15](#).

If a split brain error occurs during run time on a cluster with a direct link, perform the procedure in [“To Investigate Split Brain on Clusters With a Direct Link” on page 15](#). If a split brain error occurs during run time on a cluster without a direct link, perform the procedure in [“To Investigate Split Brain During Run Time on Clusters Without a Direct Link” on page 35](#).

▼ To Investigate Split Brain During Run Time 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 3.0 1/08 Foundation Services 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 26](#).**
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 3.0 1/08 Foundation Services 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 3.0 1/08 Foundation Services Manual Installation Guide for the Solaris OS*.

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

```
# pkill -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 17](#).

Recovering From Node Reboot at Run Time

For information about the causes of node reboot at run time, see [“A Monitored Daemon Fails Causing a Node to Reboot at Run Time”](#) on page 39.

A Monitored Daemon Fails Causing a Node to Reboot at Run Time

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 `nhpmd1M` 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 19 and [“A Monitored Daemon Fails Causing a Diskless Node or Dataless Node to Reboot at Startup”](#) on page 32.

TABLE 6-1 and TABLE 6-2 summarize the events that can cause a monitored daemon to fail at run time. To recover from daemon failure, perform the procedure in “To Recover From Daemon Failure” on page 40.

TABLE 6-1 Possible Causes of Daemon Failure at Run Time

Failed Daemon	Possible Cause at Run Time
nhcmmd	The nhcmmd daemon was killed. The failing node does not see its presence in the cluster_nodes_table.
nhprobed	The nhprobed daemon was killed.

TABLE 6-2 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.
nhcmmd	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 on the Solaris OS and the `/var/sun/nhas/core` directory on Linux, and has the format: `core.node_name.executable_file_name.process_ID.time`

For more information about core dumps, see the `coreadm1M` man page on Solaris or `coreadm8` on Linux.

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 `nhpmd1M` man page.

3. Identify the cause of the daemon failure.

Use the information obtained in [Step 1](#), [Step 2](#), [TABLE 6-1](#), and [TABLE 6-2](#).

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 3.0 1/08 Foundation Services 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 3.0 1/08 Foundation Services Getting Started Guide*. For information about how to add a node to a running cluster, see the *Netra High Availability Suite 3.0 1/08 Foundation Services 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 3.0 1/08 Foundation Services Cluster Administration Guide*. If the node or nodes are not added successfully, perform the following procedure.

▼ 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 3.0 1/08 Foundation Services Getting Started 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 `nhcmmstat1M` 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_9` or `Solaris_10`.

- 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 3.0 1/08 Foundation Services Manual Installation Guide for the Solaris OS*.

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 08:00:20:f9:b3:aa 08:00:20:f9:b3:ab node30
NODE=40 08:00:20:f9:b3:bb 08:00:20:f9:b3:bc 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.conf4man` 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 47](#)
- [“NMA Not Restarted After Failure” on page 48](#)
- [“NMA Not Sending SNMP Traps to a Given Target” on page 49](#)
- [“The `switchOver` Method Does Not Finish Executing” on page 49](#)
- [“Cascading Fails” on page 50](#)

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 48.

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 3.0 1/08 Foundation Services 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 3.0 1/08 Foundation Services 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 3.0 1/08 Foundation Services Reference Manual*. For information about sending SNMP traps, see “Developing a Remote SNMP Manager” in the *Netra High Availability Suite 3.0 1/08 Foundation Services 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 48.

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 3.0 1/08 Foundation Services 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 48](#).

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 Netra HA Suite 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 run time, 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 Netra HA Suite, see the following sections:

- [“Introduction to Error Messages” on page 53](#)
- [“Error Messages Written During Installation” on page 54](#)
- [“Error Messages Written During Run Time” on page 56](#)
- [“Error Messages Written by Command-Line Tools” on page 65](#)

Introduction to Error Messages

All of the Netra HA Suite 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 3.0 1/08 Foundation Services 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 Netra HA Suite.

All error messages have the following format:

<code><date><nodeid><PID><stack><failed service or daemon><message></code>
--

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 Run Time

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 3.0 1/08 Foundation Services 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 3.0 1/08 Foundation Services Reference Manual*.

```
Another CMM is running => exit
```

A second instance of the `nhcmmd` 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_table4` man page.

[Config] Could not initialize Cluster nodes table (NFS)

Confirm that the `cluster_nodes_table` file is correctly configured. See the `cluster_nodes_table4` man page.

[Config] Could not initialize minimal config

Confirm that the `target.conf` file is correctly configured. See the `target.conf4` 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 3.0 1/08 Foundation Services 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 3.0 1/08 Foundation Services Reference Manual*.

[Config] Invalid node name in file %s

Confirm that the `cluster_nodes_table` file is correctly configured. See the `cluster_nodes_table4` 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 3](#).

[Config] Nodes number exceed cluster capacities

For information about the supported cluster configuration, see the *Netra High Availability Suite 3.0 1/08 Foundation Services Getting Started 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 17.

[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:

1. Boot the master-eligible nodes in single-user mode:
2. Reset the replication configuration on both nodes:
3. Recreate an empty replicated configuration file by typing Y at on both nodes:
4. Reboot the nodes.
5. If the problem persists:
 - Boot both master-eligible nodes in single user mode.
 - On each master-eligible node, edit the `/etc/opt/SUNWcgsha/target.conf` file by setting the attributes field to “—” and the role field to “—”.
For information about the `target.conf` file, see the `target.conf4` man page on the Solaris OS or the `target.conf5` man page on Linux.
 - 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 3.0 1/08 Foundation Services Manual Installation Guide for the Solaris OS*.

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.conf4` 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 `nhcmmmd` 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 3.0 1/08 Foundation Services 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 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 nhcmmnd 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 Netra HA Suite command-line tools. For information about these tools, see their man pages in section 1M for the Solaris OS or 8 for *Linux of the Netra High Availability Suite 3.0 1/08 Foundation Services 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 Netra HA Suite 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 node name 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 config 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 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 3](#) and [Chapter 4](#).

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 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, 43

B

boot

- continuous reboot cycle, 29

- dataless node does not boot at startup, 31

- diskless node does not boot

 - after failover, 36

 - at startup, 25

- master-eligible node does not boot at startup, 11

C

cannot resolve a problem, 3

cascading fails, 50

Cisco switches, confirming that the spanning tree protocol is disabled, 26

cluster

- adding nodes, 43

- cannot add nodes by using nhinstall, 43

- confirming Solaris environment for diskless nodes, 44

Cluster Membership Manager, error messages, 56

command-line tools, error messages, 65

commands for debugging, 3

communication, testing between master-eligible nodes, 16, 36

configuration files for debugging, 3

continuous reboot cycle, 29

customer support center, when to contact, 3

D

Daemon Monitor

- cannot restart the NMA, 48

- finding and resetting the reset retry count, 48

daemons, causes of failure at runtime, 40

dataless node

- does not boot at startup, 31

- reboots because a monitored daemon fails

 - at runtime, 39

 - at startup, 32

- Solaris Operating System does not start, 31

DHCP

- configuration, 26

- container files, 27

diskless node

- debug using snoop, 25

- does not boot

 - after failover, 36

 - at startup, 25

- Foundation Services do not start, 29

- reboots because a monitored daemon fails

 - at runtime, 39

 - at startup, 32

- Solaris environment for, 44

- Solaris Operating System does not start, 26

E

election

- no master node at startup, 13

- two master nodes

 - at runtime, 35

 - at startup, 15

endless reboot, recovery from, 60

error messages

- Cluster Membership Manager, 56

- command-line tools, 65

- format, 53

- installation, 54

- introduction to, 53

- manual installation, 54

- nhadm, 66

- nhcrfsadm, 71

- nhinstall, 54

- NMA, 64

- Reliable Boot Service, 64

- Reliable NFS, 60

- runtime, 56

F

failover

- diskless node boot fails, 36

- two master nodes, 35

FAQs, 1

floating external address used by the NMA, 49

Foundation Services do not start, 29

H

HTTP clients, finding maximum number, 50

I

installation

- error messages, 54

- incorrect software installed, 5

- nhinstall problems, 7

- problems, 5

- stops while using nhinstall, 7

J

Java Dynamic Management Kit connector

- allocating a server port number, 21

M

manual installation error messages, 54

master node

- cannot synchronize at startup, 17

- no master node at startup, 13

- two master nodes

 - at runtime, 35

 - at startup, 15

master-eligible node

- does not boot at startup, 11

- reboots because a monitored daemon fails

 - at runtime, 39

 - at startup, 19

- Solaris Operating System does not start, 11

- startup problems, 11

- testing communication between master-eligible

 - nodes, 16, 36

master-ineligible node

- reboots because a monitored daemon fails

 - at runtime, 39

 - at startup, 32

monitored daemon

- causes of failure at runtime, 40

- fails at runtime, 39

- fails at startup

 - of diskless and dataless nodes, 32

 - of master-eligible node, 19

mount point, verifying, 19

N

nhadm check starting, testing communication

- between master-eligible nodes, 16, 36

nhadm error messages, 66

nhcmmstat, finding

- cluster configuration, 43

- role of a node, 16, 35

nhcrfsadm error messages, 71

nhcrfsd, verifying operation of, 18

nhinstall

- adding nodes to a cluster, 44

- cannot add nodes to a cluster, 43

- error messages, 54

- problems, 7

- stops during installation, 7

NMA

- allocating a server port number for the Java
Dynamic Management Kit connector, 21

- cascading fails, 50

- does not

 - communicate with external client, 47

 - initiate switchover, 49

 - restart after failure, 48

 - send SNMP traps, 49

- error messages, 64

- exists at startup of master-eligible nodes, 20

- finding the maximum number of HTTP clients, 50
- making accessible through RMI, 48
- runtime problems, 47
- using the floating external address, 49
- node
 - adding to a cluster, 43
 - cannot add to a cluster by using `nhinstall`, 43
 - reboot at runtime, 39
 - reboot at startup
 - of diskless and dataless nodes, 32
 - of master-eligible node, 19
 - Solaris environment for diskless nodes, 44

Q

- questions commonly asked, 1

R

- reboot
 - at runtime, 39
 - at startup
 - of diskless and dataless nodes, 32
 - of master-eligible node, 19
 - continuous reboot cycle, 29
 - diskless node does not reboot after failover, 36
- Reliable Boot Service error messages, 64
- Reliable NFS
 - error messages, 60
 - verifying that the `nhcrfsd` daemon is running, 18
- reset retry count, finding and resetting, 48
- RNFS.Share property, entry in `nhfs.conf`, 19
- role of a node, finding, 16, 35, 43
- runtime
 - causes of node reboot, 39
 - error messages, 56
 - monitored daemon fails, 39
 - NMA problems, 47
 - problems with failover and switchover, 35
 - two master nodes, 35

S

- SNMP
 - configuration files, 49
 - traps not being sent by the NMA, 49
- snoop, 25
- software, incorrect software installed, 5
- Solaris environment for diskless nodes, 44

- Solaris Operating System does not start
 - on a dataless node, 31
 - on a diskless node, 26
 - on a master-eligible node, 11
- spanning tree protocol, confirming disabled, 26
- split brain
 - at runtime, 35
 - at startup, 15
- startup
 - cannot synchronize master node and vice-master node, 17
 - dataless node does not boot, 31
 - diskless node does not boot, 25
 - Foundation Services do not start on a diskless node, 29
 - master-eligible node does not boot, 11
 - monitored daemon fails
 - on diskless and dataless nodes, 32
 - on master-eligible node, 19
 - NMA exists at startup of master-eligible nodes, 20
 - no master node, 13
 - of master-eligible nodes, 11
 - problems with installation, 5
 - Solaris Operating System does not start
 - on a dataless node, 31
 - on a diskless node, 26
 - on a master-eligible node, 11
 - two master nodes, 15
- switchover
 - cannot be initiated by the NMA, 49
 - two master nodes, 35
- switchOver method does not finish executing, 49
- synchronization
 - cannot synchronize at startup, 17
 - verifying mount point, 19

T

- tools for debugging, 3

V

- vice-master node
 - cannot synchronize at startup, 17
- two master nodes
 - at runtime, 35
 - at startup, 15

