



Netra™ High Availability (HA) Suite Foundation Services 2.1 7/05 Release Notes

Sun Microsystems, Inc.
www.sun.com

Part No. 817-2333-10
April 2006, Revision 03

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

Copyright 2005 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 in other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, OpenBoot, SunVTS, SunSolve, Netra, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and in other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and in 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.

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 2005 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, Californie 95054, Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. a les droits de propriété intellectuels relatants à la technologie qui est décrit dans ce document. En particulier, et sans la limitation, ces droits de propriété intellectuels 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, the Sun logo, AnswerBook2, docs.sun.com, OpenBoot, SunVTS, SunSolve, Netra, 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 AUTORISEE 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.

Contents

Introduction	2
What's New	2
New Functionalities	2
Data Sharing Between Master Nodes Using Shared Disks	3
IPv6 and IPMP Support on External Addresses	3
CGTP and IPMP Sharing a Link Over VLANs	3
Failover Resulting From Master External Link Failure	4
32-Node Support	4
SAF/CLM API Support	4
Rolling Upgrade Functionality	5
Product and Support Changes	5
SMCT Deprecation	6
CMM API Library Size	6
Node State Management Component Replacement	6
Sun StorEdge Network Data Replicator (SNDR) Software Renamed Sun AVS Software	6
Installation	7
The <code>nhinstall</code> Tool	7
Automating Patch Installation	7
Assigning IP Addresses With Diskless Nodes and DHCP	7

Manual Installation	8
Supported Hardware	8
Supported Software Versions	9
Supported Software	9
Embedded Software	10
Development Tools	10
Software Platform Version	10
Software Patches	10
Solaris OS Patches for IPMP	11
SNDR Patch	11
Carrier Grade Transport Protocol (CGTP) Patches	11
Product Recommendations	12
Use of the Reboot Command	12
Scheduling Major Tasks When the Cluster Is Unsynchronized	12
Known Issues	12
Cluster Membership Manager (CMM)	13
Reliable NFS Known Issues	16
CGTP Known Issues	17
Reliable Boot Service (RBS) Known Issues	18
Documentation Errata and Addenda	20
Setting the <code>auto-boot-retry</code> Variable	20
SMCT Deprecation	20
SAF/CLM Support	20
Update to the Custom Installation Guide	21
Consolidation of the What's New Guide	21

Netra High Availability Suite Foundation Services 2.1 7/05 Release Notes

The *Netra High Availability (HA) Suite Foundation Services 2.1 7/05 Release Notes* contain important and late-breaking information about the current release of the Foundation Services product. This document contains the following sections:

- [“Introduction” on page 2](#)
- [“What’s New” on page 2](#)
- [“Installation” on page 7](#)
- [“Supported Hardware” on page 8](#)
- [“Supported Software Versions” on page 9](#)
- [“Software Patches” on page 10](#)
- [“Product Recommendations” on page 12](#)
- [“Known Issues” on page 12](#)
- [“Documentation Errata and Addenda” on page 20](#)

Introduction

These release notes contain important product notes and known restrictions in the Netra™ HA Suite Foundation Services 2.1 7/05. Workarounds to known bugs are provided where possible. In cases where there are differences between these release notes and the Netra HA Suite Foundation Services 2.1 7/05 documentation set, the information in these release notes takes precedence. In the rest of this document, the product is referred to as the Foundation Services.

For information about supported hardware, see the *Netra High Availability Suite Foundation Services 2.1 7/05 Hardware Guide*. For information about the packages and patches delivered in the software delivery, see the *Netra High Availability Suite Foundation Services 2.1 7/05 README*.

If you are planning to upgrade your cluster from Foundation Services 2.1 6/03 to Foundation Services 2.1 7/05, install the documentation packages as described in the *Netra High Availability Suite Foundation Services 2.1 7/05 README*. After you have installed the documentation, see “Upgrading the Cluster” in the *Netra High Availability Suite Foundation Services 2.1 7/05 Custom Installation Guide*.

What’s New

This section describes new services, service enhancements, and product changes to the Netra HA Suite Foundation Services 2.1 7/05 software. It provides information about the following topics:

- [“New Functionalities” on page 2](#)
- [“Product and Support Changes” on page 5](#)

For further information about the services available in the Foundation Services 2.1 7/05 product, see the *Netra High Availability Suite Foundation Services 2.1 7/05 Overview*. For detailed installation information, see the *Netra High Availability Suite Foundation Services 2.1 7/05 Custom Installation Guide*.

New Functionalities

The following new functionalities are provided for Foundation Services 2.1 7/05:

- [“Data Sharing Between Master Nodes Using Shared Disks” on page 3](#)
- [“IPv6 and IPMP Support on External Addresses” on page 3](#)

- [“CGTP and IPMP Sharing a Link Over VLANs” on page 3](#)
- [“Failover Resulting From Master External Link Failure” on page 4](#)
- [“32-Node Support” on page 4](#)
- [“SAF/CLM API Support” on page 4](#)
- [“Rolling Upgrade Functionality” on page 5](#)

Data Sharing Between Master Nodes Using Shared Disks

The shared disk functionality is a newly supported method for sharing data between master nodes in a cluster. This is an alternative to the data replication over IP that has, to date, been the only method supported for use with the Foundation Services. In the Foundation Services 2.1 7/05 product, both methods are supported to provide reliable services such as Reliable NFS (RNFS) and reliable boot services (RBS). These two methods are to be used exclusively.

Note – The shared disk functionality is only supported for use on the Solaris 9 9/04 OS and forward.

IPv6 and IPMP Support on External Addresses

External addressing is now managed by the External Address Manager (EAM), which behaves much the same as the Node State Manager (NSM), but with the extended capability to support IPv6 and IPMP on external links. You can configure an IPv6 interface the same way you would configure an IPv4 address.

The IPMP capability increases cluster availability by allowing a master-eligible node (MEN) to have multiple Ethernet connections. If one connection fails, IPMP switches to the next. This takes less time than requiring a failover from the master node to the vice-master.

CGTP and IPMP Sharing a Link Over VLANs

Configurations where CGTP and IPMP share a link over VLANs are now supported for use with the Foundation Services software. Obviously, such configurations depend on VLAN support provided by network interface adapters and involved switches.

Failover Resulting From Master External Link Failure

EAM has the ability to trigger a failover if all external links to a master node are down. EAM uses the `nheadm` daemon, which is tracked by the Daemon Monitor. This fixes a previous problem where the health of the link was not monitored, so that a node—even a master node—might be working correctly, but be unreachable.

32-Node Support

Like the previous version of the Foundation Services, the current update of the product supports cluster configurations of two MENs and a number of master-ineligible nodes.

This release supports 32-node dataless clusters. Cluster performance (for example, the time required for switchover, failover, and boot) depend on the number of master-ineligible nodes in the cluster. When there are more than 18 master-ineligible nodes, we recommend that you use MENs that are more powerful than your master ineligible nodes to get expected performance results.

Note – The objective of the Foundation Services 2.1 7/05 release is to support 64-node dataless clusters. This configuration will be qualified when the hardware is available to do so.

For a full list of supported hardware and example cluster configurations, see the *Netra High Availability Suite Foundation Services 2.1 7/05 Hardware Guide*.

SAF/CLM API Support

The Foundation Services 2.1 7/05 is now supported for use with the Service Availability Forum (SAF) Cluster Membership (CLM) API, in addition to the existing CMM API. Either API can provide membership information about the nodes in a cluster. More information about the SAF CLM API can be found in the *Netra High Availability Suite Foundation Services 2.1 7/05 SAF Programming Guide* and at:

<http://www.saforum.org/>

Foundation Services 2.1 7/05 provides the following new functions through the SAF CLM API:

- `saClmInitialize`
- `saClmSelectionObjectGet`
- `saClmDispatch`
- `saClmFinalize`
- `saClmClusterTrackStart`
- `saClmClusterTrackStop`

- saClmClusterNodeGet
- saClmClusterNodeGetAsync

For information about these functions, go to <http://www.saforum.org/>

The following values apply to the SAF/CLM man pages when they are used with the Netra High Availability Suite Foundation Services::

TABLE 1 Changes to the SAF/CLM man pages

Location in man page	Value Inherited From the SAF Organization Man Pages	Value For Use With Netra HA Suite
SYNOPSIS section	Line that begins ???cc [flag...] file... include xxx.h	cc [flags...] file... -lSaClm include <saClm.h>
ATTRIBUTES section		SUNWnhsafclm External

Rolling Upgrade Functionality

The rolling upgrade functionality enables you to upgrade a cluster from the current version of Foundation Services that immediately preceded it without taking the entire cluster offline. Rolling upgrades can be performed with clusters that include diskless nodes, but cannot be used to upgrade from one version of the Solaris OS to another.

Note – For this release of the product, rolling upgrade can be used to upgrade only from Foundation Services 2.1 6/03 + Patch level 2 or above to Foundation Services 2.1 7/05.

For information about how to use this feature, refer to the *Netra High Availability Suite Foundation Services 2.1 7/05 Custom Installation Guide*.

Product and Support Changes

The following changes have been made to the Foundation Services software for release 2.1 7/05:

- “SMCT Deprecation” on page 6
- “CMM API Library Size” on page 6

- [“Node State Management Component Replacement”](#) on page 6
- [“Sun StorEdge Network Data Replicator \(SNDR\) Software Renamed Sun AVS Software”](#) on page 6

Note – LOMlite packages are no longer included in Netra High Availability Suite Foundation Services software distribution.

SMCT Deprecation

The SMCT tool is being deprecated and is not supported for use with version 2.1 7/05 of the Foundation Services or later.

CMM API Library Size

For Foundation Services 2.1 7/05, the CMM API is now a 64-bit library. This library provides the same API as the 32-bit library used in the Foundation Services 2.1 release.

The modification has no impact on existing 32-bit applications. However, 32-bit applications compiled with the 64-bit library might trigger a compilation warning. In this case, it might be necessary to replace `ulong` types with `uint32_t` types.

Node State Management Component Replacement

For Foundation Services 2.1 7/05, the management of the master node’s floating IP address (previously managed by the Node State Manager) is now handled by a new component called External Address Manager (EAM).

Note – If you require the installation of NSM for uses other than external access and you are using the `nhinstall` tool to install the software, set the `INSTALL_NSM` directive to `YES` in the `cluster_definition.conf` file.

Sun StorEdge Network Data Replicator (SNDR) Software Renamed Sun AVS Software

Sun StorEdge™ Network Data Replicator (SNDR) software has been renamed Sun™ Availability Suite (AVS) software. Refer to [“Supported Software Versions”](#) on page 9 for information about which versions of Sun AVS software are to be used with specific versions of the Solaris™ Operating System (Solaris OS).

Installation

This section summarizes the changes to the installation processes.

The `nhinstall` Tool

The `nhinstall` tool has been adapted to include the new services and service enhancements made to the Foundation Services 2.1 7/05. You can now use `nhinstall` to install a cluster containing both diskless and dataless nodes. The `nhinstall` tool has also been updated to work with the new hardware and operating system versions supported in this release. For information about how to install a cluster using the `nhinstall` tool, see the *Netra High Availability Suite Foundation Services 2.1 7/05 Custom Installation Guide*.

Automating Patch Installation

The `nhinstall` tool is delivered with the `addon.conf` configuration file. You can add to this file the names of packages and patches that are not delivered with the Foundation Services software delivery, but that should be installed. If this file is not configured or not present in the directory that contains the configuration files, the `nhinstall` tool assumes there are no additional patches or packages to be installed.

For information about how to use the `addon.conf` file, see the `addon.conf(4)` man page or the `addon.conf.template` file, which contains notes about modifying this file.

Assigning IP Addresses With Diskless Nodes and DHCP

It is recommended that you not use dynamic IP address assignment with Dynamic Host Configuration Protocol (DHCP). Doing so may cause the system to hang or fail when you boot the system, especially if your system's configuration contains diskless nodes with different models of hardware.

Manual Installation

All new services and service enhancements that are automatically installed by the `nhinstall` tool can also be installed manually. For information about the manual installation procedure, see the *Netra High Availability Suite Foundation Services 2.1 7/05 Custom Installation Guide*. For a complete list of packages to install, see the *Netra High Availability Suite Foundation Services 2.1 7/05 README*.

Supported Hardware

The following table summarizes the hardware supported with the Foundation Services 2.1 7/05 software as of publication of these release notes:

Servers	Netra T1 105 servers Netra T1 AC200 servers Netra T1 DC200 servers Netra 120 servers Netra 20 servers Sun Fire™ V210 servers Sun Fire V240 servers Sun Fire V440 servers Netra 240 servers Netra 440 servers Netra CT 820 servers Netra CT900 blade servers
Boards	Netra CP2300 boards with Netra CT 820 servers Netra CP2300 boards with Rapid Development Kit (RDK) Netra CP3010 ATCA SPARC blades Netra CP3020 ATCA Opteron blades
Ethernet Cards	Ethernet 10/100 1 Gbit
Disks	SCSI disks FC-AL disks IDE disks Sun StorEdge 3310 disk array

Note – Netra CT410/810 servers are not supported for use with this update of the Foundation Services software. This is because the Solaris 9 9/04 OS is the minimum update for running the Foundation Services 2.1 07/05 software, and these platforms are not supported on the Solaris 9 9/04 OS.

Supported Software Versions

This section lists the software you can use with the Foundation Services and specifies the supported versions for different types of hardware.

Supported Software

The version of Solaris software you install on a cluster depends on the hardware you use.

TABLE 2 Hardware and Recommended Solaris OS Version

Server and Boards in Use	Solaris OS Version
Netra 120, Netra CT 820, Netra T1, Netra 20, Netra CP2300 servers and board	Solaris 8 2/02 OS
Netra 240, Netra 440, Sun Fire™ V210, Sun Fire V240, Sun Fire V440 servers	Solaris 8 7/03 OS
Netra T1, Netra 20, Netra 240, Netra 440, Sun Fire V210, Sun Fire V240, Sun Fire V440, Netra CP2300 servers and board	Solaris 9 9/04 OS

For example cluster configurations, see the *Netra High Availability Suite Foundation Services 2.1 7/05 Hardware Guide*.

The following volume management software is supported for use with the Foundation Services software:

- Solstice DiskSuite™ 4.2.1 for the Solaris 8 2/02 OS. For installation information, see the *Solstice DiskSuite 4.2.1 Installation and Product Notes*.
- Solaris™ Volume Manager software for the Solaris 9 9/04 OS. For installation information, see the *Solaris Volume Manager Administration Guide*.

Embedded Software

The following software is embedded with the release of Foundation Services 7/05:

- Data replication

The following version of Sun AVS software is supported on the specified versions of the Solaris OS.

- AVS version 3.1 (with patch 116710, which is a Foundation Services-specific patch) for the Solaris 8 2/02 and 7/03 OS and Solaris 9 OS.

Note – AVS 3.2 is not supported for use with the Foundation Services software.

- Java[®] Dynamic Management Kit 5.0 software

Development Tools

The following development tools are supported for use with this release of the Foundation Services software:

- Java 2[™] Software Development Kit Standard Edition
 - Version 1.3.1 for the Solaris 8 2/02 OS and the Solaris 8 7/03 OS
 - Version 1.4.x for the Solaris 9 9/04 OS
- Sun[™] Studio 10 software

Software Platform Version

Depending on the hardware you are using, you might require a specific software platform.

- Netra[™] CT 820 servers require at least DVD0-11 when using the Solaris 8 2/02 OS
- Netra[™] CP2300 boards with the Rapid Development Kit chassis, you require at least DVD0-10 when using the Solaris 8 2/02 OS

Software Patches

The `nhinstall` tool automatically installs these patches for you. However, if you are manually installing the Foundation Services software, visit the SunSolveSM web site to download the required patches:

<http://www.sun.com/sunsolve>

Solaris OS Patches for IPMP

If you are manually installing the Foundation Services software, install the following patches that support the IPMP feature depending on the version of the Solaris OS installed on your system:

- Solaris OS 8: 108727-22 and 111958-02 (or later)
- Solaris OS 9: 115683-03 and 112911-14 (or later)

SNDR Patch

The Netra HA Suite download contains one SNDR patch: 116710-01. This SNDR/AVS point patch replaces the SNDR patches released with the previous version of the software (113054-04/, 113055-01/, and 113057-03/).

This SNDR patch is available on Sun Solve at <http://sunsolve.sun.com/point>.

Carrier Grade Transport Protocol (CGTP) Patches

The software patches for CGTP that you install depends on the version of Solaris software you are installing on the cluster. Use the following table to choose the correct patches for CGTP to install on your cluster.

TABLE 3 CGTP Patches

Solaris OS Version	Solaris Software Patch for CGTP	Location of Patch
Solaris 8 2/02 OS	112281-03	Part of Foundation Services distribution
Solaris 8 2/02 OS and kernel patch 108525-21	116036-03	Part of Foundation Services distribution
Solaris 8 OS 7/03	116036-03	Part of Foundation Services distribution
Solaris 9 9/04	No CGTP patches	

CGTP patches are point patches. They are part of the Foundation Services distribution, but they can also be downloaded from the SunSolve Web site at:

<http://sunsolve.sun.com/point>

Product Recommendations

The following sections describe recommended uses of particular functionalities and features of the Foundation Services.

Use of the Reboot Command

When rebooting a master-eligible node on a running cluster, do not use the `reboot` command. Instead, use the `init` command as root user, as follows:

```
# init 6
```

Using the `reboot` command kills processes in an indeterminate order and therefore, does not respect the required sequence for stopping services. This can lead to inconsistencies in data replication.

Scheduling Major Tasks When the Cluster Is Unsynchronized

When a master-eligible node is reintegrated into the cluster (for example, after maintenance or failure), there is a period when disk partitions are resynchronizing. While a cluster is unsynchronized, the data on the master node disk is not fully backed up. Do not schedule major tasks when the cluster is unsynchronized.

Known Issues

This section lists the known bugs and their workarounds where available.

Cluster Membership Manager (CMM)

TABLE 4 Known Issues for CMM

Bug	Description
4697437	<p>Notifications of Diskless Node State Transitions Can Be Lost</p> <p>Notifications that describe the difference between an initial state and a final state are emitted by the CMM on the master node when the cluster membership changes. The CMM running on a diskless node can miss notifications for transitory states. For example, when a cluster passes through three states (CC1, CC2, and CC3), a notification should be emitted to describe the transition from CC1 to CC2, and then to describe transition from CC2 to CC3. In this release of the product, a diskless node might only receive the notification for the overall transition from CC1 to CC3. The diskless node might miss the notification for the transient state CC2.</p> <p>When a cluster passes from state CC1 to CC2, and then back to state CC1, the diskless node might not receive any notification.</p>
4746183	<p>Single Point of Failure Occurs Immediately After Switchover</p> <p>A single point of failure exists for a brief period of time after a switchover. The single point of failure lasts until the Reliable NFS receives the following notifications from the CMM, MASTER_ELECTED and VICE_MASTER_ELECTED.</p> <p>You can use the <code>nhcmmstat</code> tool to check which notifications have been received.</p> <p>If the newly elected master node reboots before the notifications are received, refer to the <i>Netra High Availability Suite Foundation Services 2.1 7/05 Cluster Administration Guide</i> for information about how to recover a cluster.</p>
4740446	<p>Switchover Is Initiated Even Though the CMM_FLAG_SYNCHRO_NEEDED Flag Is Set</p> <p>There is a small time frame between the issuance of a command to change the synchronization state at the API level and the moment when the <code>nhcmmnd</code> daemon handles the command. If a switchover request is issued within the time, the switchover request is accepted even if the cluster is no longer synchronized.</p> <p>In this scenario, a call from Reliable NFS to clear the <code>CMM_FLAG_SYNCHRO_NEEDED</code> flag will fail because a switchover is in progress. Therefore, the master node reboots and the replication stops until the vice-master node is rebooted.</p> <p>Verify that the <code>CMM_FLAG_SYNCHRO_NEEDED</code> flag is clear before requesting a switchover. To recover from this problem, reboot the vice-master node.</p>

TABLE 4 Known Issues for CMM (Continued)

Bug	Description
4751051	<p data-bbox="476 241 1193 263">Heartbeat of the Master Node Can Be Lost During Synchronization</p> <p data-bbox="476 274 1222 439">During a full synchronization between the master node and the vice-master node, the following message might appear on the vice-master node console: master loss detected, but cannot switchover. This message is generated because the network load prevents the vice-master node from detecting all of the master node heartbeats. The vice-master node can, therefore, conclude that the master node has failed.</p> <p data-bbox="476 449 1222 499">Because the synchronization is in progress, the vice-master node cannot take the master role and there is no impact on the master node.</p> <p data-bbox="476 510 1222 562">During periods when the vice-master node cannot detect the heartbeat of the master node, the synchronization is paused.</p>
4749139	<p data-bbox="476 583 1076 605">Library Clients Should Rely on Local Notifications Only</p> <p data-bbox="476 616 1222 696">When a master-eligible node is elected as the vice-master node, the master node notifies the other peer nodes just before the data in the master node API module is updated.</p> <p data-bbox="476 706 1222 812">As a result, the <code>cmm_vicemaster_getinfo()</code> function called on the master node can fail and return a <code>CMM_ESRCH</code> error, even though the CMM library clients on the other peer nodes have already received the <code>CMM_VICEMASTER_ELECTED</code> notification.</p> <p data-bbox="476 822 1222 874">See the <i>Netra High Availability Suite Foundation Services 2.1 7/05 CMM Programming Guide</i> for more information.</p>
4796226	<p data-bbox="476 895 1222 946"><code>cmm_mastership_release</code> on Master Node Returns Incorrect Value if Vice-Master Out</p> <p data-bbox="476 956 1222 1062">When the <code>cmm_mastership_release</code> function is run on the master node, the function checks for the presence of the vice-master node. If the vice-master node is <code>OUT_OF_CLUSTER</code>, the function should return the <code>CMM_ECANCELED</code> value.</p> <p data-bbox="476 1072 1222 1102">Instead, the function returns <code>CMM_ETIMEDOUT</code> value.</p>

TABLE 4 Known Issues for CMM (Continued)

Bug	Description
4854761	<p data-bbox="555 241 1253 291">cmm_membership_remove on Master Node Causes Errors if Vice-Master Node Fails</p> <p data-bbox="555 302 1299 380">If the vice-master node fails while the <code>cmm_membership_remove</code> function is running on the master node, the <code>CMM_OK</code> value is returned but the master node does not behave correctly.</p> <p data-bbox="555 390 933 413">The master node does the following:</p> <ul data-bbox="555 423 1172 517" style="list-style-type: none"><li data-bbox="555 423 915 446">• Continues to be the master node<li data-bbox="555 456 1115 479">• Does not detect that the vice-master node has failed<li data-bbox="555 489 1172 517">• Does not detect any other changes in cluster membership
4845598	<p data-bbox="555 536 1282 586">Diskless Node Emits CMM_INVALID_CLUSTER Notification When Master Is Disqualified</p> <p data-bbox="555 597 1299 708">When the master node is disqualified by the <code>cmm_membership_qualif</code> function, the <code>nhcmmmd</code> daemon on an associated diskless node might emit a <code>CMM_INVALID_CLUSTER</code> notification. Ignore the notification. The cluster is up and running.</p>
4928087	<p data-bbox="555 727 1233 777">switchover + full synch operation generates a duplicate floating address</p> <p data-bbox="555 788 1299 899">When you <code>switchover (/opt/SUNWcgha/sbin/nhcmmstat -c so)</code> in parallel with a full synchronization when the two master-eligible nodes are synchronized (<code>/opt/SUNWcgha/sbin/nhrfsadm -f</code>), the following events occur:</p> <ul data-bbox="555 909 1293 1083" style="list-style-type: none"><li data-bbox="555 909 1043 932">• The <code>nhcmmmd</code> daemon engages the switchover<li data-bbox="555 942 1253 992">• The starting of a full synchronization of the master-eligible nodes changes the nodes' state from <code>READY</code> to <code>SYNCHO NEEDED</code><li data-bbox="555 1003 1293 1083">• The vice-master becomes master and sets its master IP address to <code>UP</code> (The master IP address is always plumbed on both master-eligible nodes but this address is set to <code>DOWN</code> on the vice-master node). <p data-bbox="555 1093 1299 1171">As a result of this sequence of events, the Reliable NFS cannot set the master IP address to <code>DOWN</code> because this action cannot take place while a full synchronization is in progress.</p> <p data-bbox="555 1182 1253 1236">If you encounter this problem, wait until the full synchronization is complete. This might take some time.</p>

Reliable NFS Known Issues

TABLE 5 Known Issues for Reliable NFS

Bug	Description
4624575	<p>Clients Hang When the Vice-Master Node is Stopped</p> <p>Halting the vice-master node when clients are writing data to the master node might cause clients to hang for up to 16 seconds before continuing processing.</p> <p>This problem does not occur when the vice-master node is shut down using the procedures described in the <i>Netra High Availability Suite Foundation Services 2.1 7/05 Cluster Administration Guide</i>.</p>
4960188	<p>NFS client server deadlock</p> <p>The Solaris OS does not support the ability for a node with an NFS client to access data exported by an NFS server on the same host. In this case, if the NFS client writes large files to the NFS server, the OS deadlocks and the node hangs.</p> <p>This situation can occur if an application on the vice-master fails over or is switched over to the master. The master node hangs and might not be able to function as a master. If you encounter this situation, reboot the hung node.</p>
4964345	<p>SNDR sets using sector 0 fail, which is not detected by nhinstall/nhadm</p> <p>Do not use sector zero in any slice that will be replicated. If you do, the cluster will hang on the final step of SNDR synchronization. You might encounter this situation if you install more than one disk on a node using nhinstall.</p>

CGTP Known Issues

TABLE 6 Known Issues for CGTP

Bug	Description
4740370	<p data-bbox="555 326 1250 347">CGTP Broadcast IRE Are Not Recreated After plumb or unplumb</p> <p data-bbox="555 357 1279 439">Use of the <code>ifconfig</code> command to plumb or unplumb the CGTP interface is not supported. Using the <code>ifconfig</code> command in this way can lead to unexpected cluster outage.</p> <p data-bbox="555 449 1300 586">Action on a single interface leads to inoperative CGTP broadcasts. Broadcasts replicated by CGTP might not be delivered if one of the underlying incoming interfaces is down, and, for the same reason, if the interface has been unplumbed. CGTP broadcasts can NOT survive the brutal unplumbing/replumbing of the underlying network interfaces. The only way for CGTP broadcasts to survive an <code>ifconfig</code> unplumb is to always respect the following sequence of operations:</p> <ul data-bbox="555 656 1258 777" style="list-style-type: none"><li data-bbox="555 656 1258 677">• Delete the CGTP routes that cross the interface being unplumbed.<li data-bbox="555 687 818 708">• Unplumb the interface.<li data-bbox="555 718 843 739">• Replumb a new interface.<li data-bbox="555 749 953 770">• Redeclare the previous CGTP routes

Reliable Boot Service (RBS) Known Issues

TABLE 7 Known Issues for RBS

Bug	Description
4621703	<p>Boot Server Allocates the Same IP Address to Two Diskless Nodes</p> <p>When a diskless node is either stopped or disconnected from the network and then restarted or reconnected without rebooting the OS, the first diskless node's IP address can be allocated to another diskless node that is booted at the same time. This can happen in any of the following situations:</p> <ul style="list-style-type: none">• Diskless nodes are configured using the DHCP dynamic boot policy.• One diskless node is stopped or disconnected through both Ethernet links for some time and another diskless node boots during this time.• There are no more IP addresses available in the pool of DHCP addresses managed by the RBS.• The stopped/disconnected node is restarted/reconnected without rebooting the operating system. Consequently, a booting diskless node could be allocated the IP address of the stopped/disconnected node. When restarted/reconnected, the first node would keep the IP address. This scenario would cause network problems for both nodes. To recover, reboot the node that has been stopped/disconnected.
6207722	<p>path_to_inst corruption</p> <p>This can occur if the host is suspended by a console break, then booted. Its root cause is bug 4520944, which is fixed in the Solaris 9 OS by patch 116548-03.</p>
6208336	<p>CMM_ETIMEDOUT errors is displayed when performing a switchover on 64 nodes</p> <p>On large clusters (for example, 40 plus-node clusters), this error appears on the master node when you perform a switchover using <code>nhcmmstat</code> command, even if the switchover succeeds.</p> <p>To resolve the error described above, Foundation Services 2.1 7/05 now allows you to define the timeout period (in seconds) for a particular run of the <code>nhcmmstat</code> command. This is done using the <code>-m</code> option of the <code>nhcmmstat</code> command. If you receive the above error, increase the value of this option.</p> <p>The command line argument is <code>-m <timeout></code>. The default value for this option is five seconds. The following example shows how to trigger a switch-over with a time-out of 6 seconds.</p> <pre>/opt/SUNWcgha/sbin/nhcmmstat -m 6 -c so</pre>

TABLE 7 Known Issues for RBS

Bug	Description
6218803	<p>DHCP table corruption</p> <p>Each diskless node has its own <code>dhcpageant</code> file in the exported / (root) partition on the master server. For example, <code>/export/root/[nodename]/etc/default/dhcpageant</code></p> <p>It is possible for this file to become corrupted when the diskless node crashes or goes down without a filesystem sync. If this occurs, the diskless node will not boot until the file is repaired on master server.</p> <p>To avoid file corruption, you can keep local copies of such dhcp tables on master and vice-master servers. The synchronization of these local copies is then out of scope from NHAS Foundation Service software and falls under manual system administration.</p> <p>You can use <code>REPLICATED_DHCP_FILES=NO</code> option in the <code>cluster_definition.conf</code> file while using <code>nhinstall</code>.</p>
6267056	<p>Only ViceMaster Node will stay up in the cluster when performing switchover+full synchronization</p> <p>When you switchover (<code>/opt/SUNWcgha/sbin/nhcmstat -c so</code>) there is a brief window during which launching a full synchronization (<code>/opt/SUNWcgha/sbin/nhcrfsadm -f</code>) will cause the cluster to loose its master node (and, therefore, the diskless nodes) and only the vice master node will stay up.</p> <p>If you encounter this problem, recover the cluster by flushing the SNDR configuration. For more information about recovering a cluster, see the <i>Netra High Availability Suite Foundation Services 2.1 7/05 Troubleshooting Guide</i>.</p>
6290647	<p>On large clusters, when one of the MEN joins back the cluster, users might see some warnings on the console of the joining node: <code>/var/run/CMM_xxx_00000000 fails: Resource temporarily unavailable</code></p> <p>This means that some membership/mastership notifications have been lost in that node. You might want to have client applications update their view of the cluster to a coherent state by calling the CMM API <code>cmm_member_getall()</code> function or the SAF CLM API <code>saClmClusterTrack()</code> function with the <code>SA_TRACK_CURRENT</code> flag.</p>
6324905	<p>The problem described here is not related to RBS, but occurs during reboot. If your system has an incorrect system date, clusters will reboot in a loop, returning the error, "failed to open <code>getexecname()</code>". This issue might occur on Netra CT 820 servers when a board is removed or replaced.</p> <p>To avoid this error, and before installing and running the Netra HA Suite Foundation Services software on any platform, ensure that all nodes have the correct system date.</p>

Documentation Errata and Addenda

This section describes corrections and changes that apply to the Netra High Availability Suite Foundation Services 2.1 7/05 documentation set.

Setting the auto-boot-retry Variable

The *Netra High Availability Suite Foundation Services Quick Start Guide* references the `auto-boot-retry` variable. If this variable exists on your system, it must be set to true; if it does not exist on your system, disregard tasks that reference setting it.

SMCT Deprecation

Many references to the SMCT installation tool have been removed from the Foundation Services documentation set because this feature is not supported for use with Foundation Services 2.1 7/05.

After this release of the product, the *Netra High Availability Suite Foundation Services SMCT Programming Guide* will be removed from the documentation set. In addition, the following man pages, which currently have interface levels set to `obsolete` will be removed from the product download and the *Netra High Availability Suite Foundation Services 2.1 7/05 Reference Guide*:

`flinstall(1M)`, `flconfig(1M)`, `flcreate(1M)`, `fldeploy(1M)`,
`nhsmctsetup(1M)`, `slconfig(1M)`, `slcreate(1M)`, `sldelete(1M)`,
`sldeploy(1M)`, `slexport(1M)`, `cluster.conf(4)`, `install_server(4)`,
`machine.conf(4)`, `master-system.conf(4)`, `network.conf(4)`,
`software.conf(4)`, `userapp.conf(4)`

SAF/CLM Support

Due to new support of Foundation Services software for use with the SAF/CLM API, a new guide, the *Netra High Availability Suite Foundation Services 2.1 7/05 SAF Programming Guide*, has been added to the existing documentation set.

Update to the *Custom Installation Guide*

Procedures described in Chapter 9, “Upgrading the Cluster,” of the *Netra High Availability Suite Foundation Services 2.1 7/05 Custom Installation Guide* have been modified. If you are performing the tasks described in this chapter, ensure that you have downloaded the updated version of this guide, which is included in the package in which these release notes have been delivered (“Revision 02” on the cover page of the guide).

Consolidation of the *What’s New Guide*

The information previously contained in the *Netra High Availability Suite Foundation Services What’s New Guide* is now provided in the *Netra High Availability Suite Foundation Services Release Notes*.

