



# What's New in Netra High Availability Suite Foundation Services 2.1 6/03

---

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

Part No: 817-1763-11  
September 2004

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

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

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

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

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

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

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

---

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

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

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

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

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

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



040813@9495



# Contents

---

<b>Preface</b>	<b>5</b>
<b>1 What's New at a Glance</b>	<b>9</b>
Purpose of Patch Release	9
Highlights of This Delivery	10
<b>2 Installation, Configuration, and Administration</b>	<b>11</b>
Supported Cluster Configurations	11
Installation	12
Patches	12
The <code>nhinstall</code> Tool	12
Manual Installation	13
Cluster Configuration File	13
New Tools	14
<b>3 Service Enhancements</b>	<b>15</b>
Introduction to Service Enhancements	15
Reduced Duration of Disk Synchronization	16
Data Management Policy	16
Masterless Cluster	17
Sanity Check of Replicated Slices	18
Delayed Synchronization	18
Serialized Slice Synchronization	19



# Preface

---

This book describes the service enhancements provided by the September 2004 patch level of the Netra™ High Availability (HA) Suite Foundation Services 2.1 6/03 product. See the patch READMEs for details of the exact patch contents and installation method.

---

## Who Should Use This Book

This book is for system administrators or system developers who are upgrading to the Foundation Services 2.1 6/03 with the patches available at the time of publication of this book.

---

## How This Book Is Organized

This book is organized as follows:

- [Chapter 1](#) provides a quick glance at the service enhancements of the Foundation Services 2.1 6/03.
- [Chapter 2](#) describes the changes to the installation process, the configuration files, and the system administration tools.
- [Chapter 3](#) describes the service enhancements of the Foundation Services product.

---

## Related Books

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

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

---

## Accessing Sun Documentation Online

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

---

## Ordering Sun Documentation

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

---

# Typographic Conventions

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

**TABLE P-1** Typographic Conventions

Typeface or Symbol	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name%</code> you have mail.
<b>AaBbCc123</b>	What you type, contrasted with onscreen computer output	<code>machine_name%</code> <b>su</b> Password:
<i>AaBbCc123</i>	Command-line placeholder: replace with a real name or value	The command to remove a file is <code>rm filename</code> .
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. Do <i>not</i> save the file. (Emphasis sometimes appears in bold online.)

---

# Shell Prompts in Command Examples

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

**TABLE P-2** Shell Prompts

Shell	Prompt
C shell prompt	<code>machine_name%</code>
C shell superuser prompt	<code>machine_name#</code>
Bourne shell and Korn shell prompt	<code>\$</code>

**TABLE P-2** Shell Prompts (Continued)

Shell	Prompt
Bourne shell and Korn shell superuser prompt	#

## What's New at a Glance

---

The Foundation Services 2.1 6/03 are a suite of reliable services that provide the foundation for next-generation network services. The Foundation Services 2.1 6/03 can be used by network equipment providers to create a highly available, dynamically scalable cluster of distributed nodes. They can also be used to augment existing frameworks. For an introduction to the service enhancements of the Foundation Services 2.1 6/03, see the following sections:

- [“Purpose of Patch Release” on page 9](#)
- [“Highlights of This Delivery” on page 10](#)

For further information about the services available in the Foundation Services 2.1 6/03 product, see the *Netra High Availability Suite Foundation Services 2.1 6/03 Overview*.

---

## Purpose of Patch Release

The current patch level of the Foundation Services 2.1 6/03 supports an increased range of hardware platforms and an additional version of the Solaris™ operating system. Many of the services provided by the Foundation Services have been enhanced to improve the availability of applications and to increase your control of the data integrity policy.

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

---

## Highlights of This Delivery

Following are the enhancements provided in the current patch level of the Foundation Services 2.1 6/03:

- Support of clusters of more than 12 nodes. For information about the cluster size supported, see the *Netra High Availability Suite Foundation Services 2.1 6/03 Release Notes*.
- Support of Sun StorEdge 3310 disk arrays for single attachment. For a full list of supported hardware and example cluster configurations, see the *Netra High Availability Suite Foundation Services 2.1 6/03 Hardware Guide*.
- Support of clusters containing both dataless and diskless nodes. You must install these clusters manually. For further information on the installation method, see the *Netra High Availability Suite Foundation Services 2.1 6/03 Custom Installation Guide*.
- Reduction of the time taken for full synchronization.
- Data management policies: Integrity, Availability, and Adaptability.
- Qualification of Foundation Services 2.1 6/03 on new hardware platforms and new versions of the Solaris operating system. For information about the hardware and software supported, see the *Netra High Availability Suite Foundation Services 2.1 6/03 Release Notes*.
- Ability to have diskless and dataless nodes running when there is no longer a master in the cluster.
- Sanity check of the blocks on the replicated slices.
- Sequential slice synchronization.
- Ability to delay the beginning of synchronization and later trigger synchronization at a time of your choice.
- Control of the order in which `nhinstall` installs patches and packages. For further information, see “Installation” on page 12.
- Update of the books in the documentation set to reflect the changes to the product described in [Chapter 2](#) and [Chapter 3](#). The new versions of the books in the documentation set have part numbers of the format 817-xxx-20.

# Installation, Configuration, and Administration

---

For a summary of the changes to the installation process, the configuration files, and the system administration tools, see the following sections:

- “Supported Cluster Configurations” on page 11
- “Installation” on page 12
- “Cluster Configuration File” on page 13
- “New Tools” on page 14

---

## Supported Cluster Configurations

Like the previous version of the Foundation Services, the current patch level of the product supports cluster configurations of two master-eligible nodes and a number of master-ineligible nodes. The limitation of 10 master-ineligible nodes no longer exists. For information about the cluster size supported, see the *Netra High Availability Suite Foundation Services 2.1 6/03 Release Notes*. You can also now have a mix of diskless and dataless nodes in a single cluster.

The following new hardware items are now supported with the current patch level of the Foundation Services.

- Netra CT 820 server
- Netra 240
- Sun Fire™ V240
- Sun Fire V210
- Sun StorEdge 3310 disk array for single attachment

For a full list of supported hardware and example cluster configurations, see the *Netra High Availability Suite Foundation Services 2.1 6/03 Hardware Guide*.

---

# Installation

This section summarizes the changes to the packages that are installed on a cluster, and the changes to the installation processes.

## Patches

The Foundation Services 2.1 6/03 distribution contains the following patches:

- 114175-0x Patch for master-eligible nodes
- 115606-0x Patch containing enhancements to services for all cluster nodes
- 115644-0x Patch for `nhinstall` tool

When you install these patches, you must also install patch 116710. This patch is not part of the Foundation Services 2.1 6/03 distribution and is only available to Foundation Services customers. For information about the availability of this patch, ask your Sun support contact.

For a full list of the Foundation Services 2.1 6/03 packages and patches, see the *Netra High Availability Suite Foundation Services 2.1 6/03 README*.

## The `nhinstall` Tool

The `nhinstall` tool has been adapted to include the enhancements made to the Foundation Services 2.1 6/03. For information about the enhancements, see [Chapter 3](#). For information about how to install a cluster using the `nhinstall` tool, see the *Netra High Availability Suite Foundation Services 2.1 6/03 Custom Installation Guide*.

The `nhinstall` tool has been adapted to include the following enhancements:

- A choice of method by which the tool installs patches
- Ability to configure DHCP on the local file system of master-eligible nodes during diskless node installation
- Ability to use a router other than the installation server. By default the router is the installation server
- Choices on how you install the Solaris operating system:
  - Installing a different Solaris package set on diskless nodes to from the one you install on the master-eligible nodes. Both package sets must be from the same Solaris distribution. For more information, see the `nodeprof.conf(4)` and `diskless_nodeprof.conf(4)` man pages.

- Identifying the appropriate OS profile to be installed even when the version of the Solaris operating system cannot be automatically detected.
- Installing a different version of the Solaris operating system on diskless nodes to the one you install on the master-eligible nodes.

## Manual Installation

The manual installation procedure has been adapted to support the installation of diskless and dataless nodes in a single cluster.

All of the new features that are automatically installed by the `nhinstall` tool can be installed manually. For information about the manual installation procedure, see the *Netra High Availability Suite Foundation Services 2.1 6/03 Custom Installation Guide*. For a complete list of packages to install, see the *Netra High Availability Suite Foundation Services 2.1 6/03 README*.

---

## Cluster Configuration File

The Foundation Services are configured by the `nhfs.conf` file. The following new parameters have been defined in the `nhfs.conf` file:

<code>RNFS.SyncType</code>	Control of the replication mode. This impacts the time it takes for a full synchronization of the vice-master and master disks.  For further information, see <a href="#">“Reduced Duration of Disk Synchronization”</a> on page 16.
<code>Cluster.DataManagementPolicy</code>	Set the data management policies to be Integrity, Availability, or Adaptability.  For further information, see <a href="#">“Data Management Policy”</a> on page 16.
<code>CMM.MasterLoss.Detection</code>	Permit the diskless and dataless nodes to stay up even when there is no master node in the cluster.  For further information, see <a href="#">“Masterless Cluster”</a> on page 17.

<code>RNFS.CheckReplicatedSlices</code>	Continuously scan the state of replicated slices.  For further information, see <a href="#">“Sanity Check of Replicated Slices”</a> on page 18.
<code>RNFS.EnableSync</code>	Delay the start of disk synchronization.  For further information, see <a href="#">“Delayed Synchronization”</a> on page 18.
<code>RNFS.SerializeSync</code>	Synchronize the master and vice-master disks one slice at a time.  For further information, see <a href="#">“Serialized Slice Synchronization”</a> on page 19.

---

## New Tools

The Foundation Services 2.1 6/03 provides a new tool, `nhenablesync`, which enables you to trigger disk synchronization if you have delayed the start of this synchronization. For information on this tool see the `nhenablesync(1M)` man page. For information on delaying disk synchronization, see the [“Delayed Synchronization”](#) on page 18.

The Foundation Services 2.1 6/03 also provides a new subcommand of `nhadm` command, `nhadm copy`. The `copy` subcommand enables you to copy files that are held locally on the master node onto the vice-master node. Examples of files you could choose to hold locally on the master-eligible nodes are the DHCP configuration files. For more information on the `copy` subcommand, see the `nhadm(1M)` man page.

## Service Enhancements

---

For a summary of the service enhancements in the Foundation Services product, see the following sections:

- [“Introduction to Service Enhancements” on page 15](#)
- [“Reduced Duration of Disk Synchronization” on page 16](#)
- [“Data Management Policy” on page 16](#)
- [“Masterless Cluster” on page 17](#)
- [“Sanity Check of Replicated Slices” on page 18](#)
- [“Delayed Synchronization” on page 18](#)
- [“Serialized Slice Synchronization” on page 19](#)

---

## Introduction to Service Enhancements

A number of enhancements have been made to the Foundation Services to enable you to prioritize data integrity, to increase the availability of your applications, or to reduce disk and network load on the cluster. These enhancements are as follows:

Prioritizing data integrity

- [“Reduced Duration of Disk Synchronization” on page 16](#)
- The Integrity and Adaptability data management policies, as described in [“Data Management Policy” on page 16](#)

Improving availability

- The Availability data management policy, as described in [“Data Management Policy” on page 16](#)
- [“Masterless Cluster” on page 17](#)

- “Delayed Synchronization” on page 18

Reducing disk and network load

- “Serialized Slice Synchronization” on page 19

Enable or disable these service enhancements to suit your cluster needs.

---

## Reduced Duration of Disk Synchronization

This feature enables you to choose how the master and vice-master disks are synchronized. You can have synchronization either of an entire master disk or of only those blocks that contain data. By choosing the latter option, you reduce the time it takes to synchronize the master and vice-master disks.

Set this feature to have the same value on both master-eligible nodes. If you change the value of the `RNFS.SyncType` property defining this feature, reboot the cluster as described in “Shutting Down and Restarting a Cluster” in the *Netra High Availability Suite Foundation Services 2.1 6/03 Cluster Administration Guide*.

Choose the synchronization method by setting the value of the `RNFS.SyncType` parameter in the `nhfs.conf` file. If you are installing the cluster with `nhinstall`, choose the synchronization method by setting the value of `SLICE_SYNC_TYPE` in the `cluster_definition.conf` file. For more information, see the `nhfs.conf(4)` and `cluster_definition.conf(4)` man pages.

---

## Data Management Policy

There are three data management policies available in the current patch level of the Foundation Services 2.1 6/03. These policies determine how the cluster behaves when a failed vice-master reboots in a cluster that has no master node. The policy you choose depends on the availability and data-integrity requirements of your cluster.

The data management policies are the following:

Integrity	Ensures that the cluster uses the most up-to-date data. The vice-master does not take the master role but waits for the old master to return to the cluster. This is the default data management policy.
Availability	Prioritizes the availability of services running on the cluster over data integrity. The vice-master node takes the master role when there is no master in the cluster. This policy triggers a full synchronization when the old master node joins the cluster as the vice-master node. This synchronization might result in the loss of any data written to the old master while the vice-master was down.
Adaptability	Prioritizes availability only if the master and vice-master disks are synchronized. Such disk synchronization would increase the level of data integrity. The vice-master checks the disk synchronization state. If the state indicates that the master and vice-master nodes are not synchronized, the vice-master node does not take the master role but waits until the old master node returns to the cluster. If the state indicates that the master and vice-master disks are synchronized, the vice-master takes on the master role without waiting for the old master to rejoin the cluster.

Choose the data management policy by setting the value of the `Cluster.DataManagementPolicy` parameter in the `nhfs.conf` file. If you are installing the cluster with `nhinstall`, choose the data management policy by setting the value of `DATA_MGT_POLICY` in the `cluster_definition.conf` file. For more information, see the `nhfs.conf(4)` and `cluster_definition.conf(4)` man pages.

---

## Masterless Cluster

This feature permits the diskless and dataless nodes to stay up even when there is no master node in the cluster. By default, this feature is disabled and so diskless and dataless nodes reboot if there is no master node in the cluster for more than a few minutes. Enable this feature if you want to reduce the possible downtime for services running on master-ineligible nodes.

If you enable this feature, you must ensure that the diskless and dataless nodes can handle the situation where they would no longer be able to access the files exported by the master node.

Activate this feature by setting the value of the `CMM.Masterloss.Detection` parameter in the `nhfs.conf` file. If you are installing the cluster with `nhinstall`, enable this feature by setting the value of `MASTER_LOSS_DETECTION` in the `cluster_definition.conf` file. For more information, see the `nhfs.conf(4)` and `cluster_definition.conf(4)` man pages.

---

## Sanity Check of Replicated Slices

This feature enables you to continuously scan the state of replicated slices. By default, this feature is disabled and it can only be used if the master and vice-master disks are synchronized. If you do not monitor the state of the replicated slices, it is possible that the vice-master disk is corrupted and so partly or completely inaccessible.

Activate this feature by setting the value of the `RNFS.CheckReplicatedSlices` parameter in the `nhfs.conf` file. If you are installing the cluster with `nhinstall`, enable this feature by setting the value of `CHECK_REPLICATED_SLICES` in the `cluster_definition.conf` file. For more information, see the `nhfs.conf(4)` and `cluster_definition.conf(4)` man pages.

---

## Delayed Synchronization

This feature enables you to delay the start of synchronization between the master and vice-master disks. By default this feature is disabled. Delay synchronization if you do not want the synchronization task to conflict with other high CPU activities. If you enable this feature, you can trigger synchronization at a later time of your choosing using the `nhenablesync` command.

Until synchronization is triggered and completed, the vice-master is not eligible to become master and so the cluster is open to a single point of failure. By delaying synchronization you are prolonging this vulnerability.

Activate this feature by setting the value of the `RNFS.EnableSync` parameter in the `nhfs.conf` file. If you are installing the cluster with `nhinstall`, specify the synchronization method by setting the value of `SYNC_FLAG` in the `cluster_definition.conf` file. For more information, see the `nhfs.conf(4)`, `cluster_definition.conf(4)`, and `nhenablesync(1M)` man pages.

---

## Serialized Slice Synchronization

This feature enables you to synchronize the master and vice-master disks one slice at a time rather than all at once. By default this feature is disabled. Enabling this feature reduces the disk and network load. However, it limits the availability of the cluster during a certain limited time period because the vice-master cannot become master before all slices have been synchronized. During this time period, the cluster is vulnerable to a single point of failure.

Activate this feature by setting the value of the `RNFS.SerializeSync` parameter in the `nhfs.conf` file. If you are installing the cluster with `nhinstall`, specify the synchronization method by setting the value of `SERIALIZE_SYNC` in the `cluster_definition.conf` file. For more information, see the `nhfs.conf(4)`, and `cluster_definition.conf(4)` man pages.

