

Netra™ High Availability Suite 3.0 1/08 Foundation Services Glossary

Sun Microsystems, Inc.
www.sun.com

Part No. 819-5238-13
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, Netra, 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.

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 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, le logo Sun, Java, AnswerBook2, docs.sun.com, 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.



Adobe PostScript

Contents

Preface v

Glossary 1

Preface

The *Netra High Availability Suite 3.0 1/08 Foundation Services Glossary* defines the specialized terms used in the Netra™ High Availability (HA) Suite 3.0 1/08 Foundation Services documentation set.

Who Should Use This Book

This book is for system administrators who are maintaining a cluster running the Foundation Services, or for system developers who are developing applications for a cluster running the Netra HA Suite software.

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 rm <i>filename</i> .

^{*} 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
Cluster administration	<i>Netra High Availability Suite 3.0 1/08 Foundation Services Cluster Administration Guide</i>	819-5235-13

Application	Title	Part Number
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 Glossary, part number 819-5238-13.

Glossary

A

- actual priority** The priority of the Foundation Services processes compared to other system processes. The actual priority of a process is a function of its base priority and its relative priority:
- $$\text{actual priority} = \text{base priority} + \text{relative priority}$$
- See also, *base priority* and *relative priority*.
- adaptability** In this context, one of the proposed data management policies.
- See also, data management policy.
- address triplet** An addressing schema that assigns each peer node an IP address for the following interfaces:
- Physical interface of the first redundant network, *NIC0*
 - Physical interface of the second redundant network, *NIC1*
 - Virtual physical interface for CGTP, *cgtcp0*
- administrative attribute** The external management viewpoint of a node.
- amnesia** An error scenario in which a cluster restarts with a newly elected master node that is assumed to have highly available data that is synchronized with the previous master node, when, in fact, it is not. When the previous master node comes back (as new vice-master node), the latest updated data on the node will be lost.
- AMF** Availability Management Framework from Service Availability Forum (SA Forum)

automated installation	Installation using the <code>nhinstall</code> tool. See also, <code>nhinstall</code> .
availability	Refers to the probability that a service is available for use at any given time. See also, redundancy, reliability, and serviceability. Also in this context, one of the proposed data management policies. See also, data management policy.
AVS	Sun StorageTek™ Availability Suite. SNDR is part of this software suite.

B

base priority	One part of the actual priority of the Foundation Services processes: actual priority = base priority + relative priority See also actual priority and relative priority.
bitmap partition	A partition that contains the scoreboard bitmap for a replicated partition on the same disk. There is one bitmap partition per replicated partition. In SNDR, bitmap partitions are referred to as “bitmap volumes.” For a listing of all other partition types, see also, partition.
bonding	On the Linux OS, channel bonding refers to a kernel driver that aggregates multiple network links into a single link. The bonding driver sends data presented to the aggregated single link through one or more network links, depending on the bonding driver configuration and which link is up. On the Solaris OS, Solaris trunking is similar to channel bonding for selected dual/quad Ethernet cards.
bonding interface	An interface to multiple network links aggregated as a single link and managed by the bonding driver.
boot policy	The method of starting a diskless node at cluster startup. See also, DHCP client ID, DHCP dynamic, and DHCP static.

C

Carrier Grade Transport Protocol	CGTP is a mechanism added at the IP level that makes IP networks reliable. A data packet sent to a CGTP address (virtual) will be duplicated and then transmitted through two IP addresses (not virtual). Every IP address
---	--

corresponds to a network interface connected to a network link. The two network links must be two completely distinct routes (no common switch or router can be used on these links).

If data packets are lost on one link, those sent on the second link should reach the destination address without any latency. Powerful filtering mechanisms are added to the receiver site to withdraw the duplicate packets with a limited impact to performance.

cascade The transfer of statistics from the Node Management Agent (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.

CGTP Carrier Grade Transport Protocol

CGTP address The IP address of a CGTP interface. See also CGTP interface.

CGTP interface A virtual physical interface over which IP packets are transferred using the CGTP. See also, Carrier Grade Transport Protocol and virtual physical interface.

client node See master-ineligible node.

CLM Cluster Management API from Service Availability Forum (SA Forum)

cluster A group of peer nodes that are connected by a network. To run the Netra HA Suite Foundation Services, a cluster must have the following characteristics:

- A server (master-eligible) node elected as the master node
 - A server (master-eligible) node elected as the vice-master node
 - All other peer nodes are client (master-ineligible) nodes, which are either dataless or diskless nodes
 - Data replication over IP or shared disks between the master node and vice-master node
 - A redundant network between all peer nodes
- See also, client node, dataless node, diskless node, master node, master-eligible node, master-ineligible node, peer node, and vice-master node.

clusterid The number that identifies a cluster. Also called domainid. The clusterid is the same for each peer node in a cluster. See also nodeid.

Cluster Membership

Manager Cluster Membership Manager is a service that manages the cluster membership and mastership. The CMM performs the following tasks:

- Determines which nodes are members of the cluster
- Assigns the roles and attributes of nodes
- Detects the failure of nodes
- Notifies clients of changes to the cluster

cluster network	The private network over which peer nodes communicate with each other. See also external network.
cluster_nodes_table	A configuration file that stores the membership and configuration information for all peer nodes in a cluster. A node must have an entry in the <code>cluster_nodes_table</code> file to be part of a cluster. There is a <code>cluster_nodes_table</code> on both master-eligible nodes in the cluster.
CMM	Cluster Membership Manager
CompactPCI	Compact Peripheral Component Interconnect
CompactPSB	CompactPCI Packet Switched Backplane
control domain	domain that creates and manages other logical domains and services.
custom installation	Manual installation that does not use the <code>nhinstall</code> tool. See also, <code>nhinstall</code> .

D

dataless node	A peer node that boots from its local disk. A dataless node runs customer applications locally, but stores and accesses reliable (highly available) data through the cluster network on a replicated or shared partition of the master node, using the Reliable File Service of the Foundation Services. A dataless node cannot be the master node or the vice-master node. For a listing of all other node types, see also, <code>node</code> .
data integrity	Keeping data correct, complete, whole, and consistent with the intent of the data's creators. It is achieved by preventing accidental or unauthorized insertion, modification or destruction of data in a database.
data management policy	The data management policy of a cluster determines 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 three data management policies are integrity, availability, and adaptability. The integrity policy is chosen when data integrity must be privileged; availability is chosen when data availability must be privileged (however, note that some data can be lost); adaptability is chosen when you are relying on the state of synchronization between the two master nodes to choose the behavior of the vice-master node.

data partition	<p>A partition that contains data. If data contained in the partition must be highly available, the data partition must be either a replicated partition (if the cluster is configured to use data replication over IP) or a shared partition (if the cluster is configured to use shared disks).</p> <p>See also IP replication and shared disk. For a listing of all other partition types, see also, partition.</p>
development server	<p>Optional hardware used to develop applications using the Netra HA Suite APIs. This hardware can run the Solaris OS or Linux distributions. Refer to the <i>Netra High Availability Suite 3.0 1/08 Foundation Services Getting Started Guide</i> for more information about using a development server.</p>
DHCP	<p>Dynamic Host Configuration Protocol</p>
DHCP client ID	<p>A boot policy that associates a <code>CLIENT_ID</code> string with a diskless node. See also, DHCP dynamic and DHCP static.</p>
DHCP dynamic	<p>A boot policy that creates a dynamic map between the Ethernet address of a diskless node and an IP address taken from a pool of available addresses. See also, DHCP client ID and DHCP static.</p> <p>DHCP dynamic is not recommended to be used with Reliable Boot Service, as it is no longer supported for this type of Reliable Service.</p>
DHCP static	<p>A boot policy that maps the Ethernet address of a diskless node to a fixed IP address. See also, DHCP client ID and DHCP dynamic.</p>
direct link	<p>A link between the serial ports on the master-eligible nodes. The direct link helps prevent the occurrence of split brain when the network between the master node and vice-master node fails.</p>
diskfull node	<p>A peer node that contains at least one disk on which applications can run and information can be permanently stored. Master-eligible nodes are diskfull nodes. Dataless nodes are not considered to be diskfull nodes. For a listing of all other node types, see also, node.</p>
diskless node	<p>A peer node that does not have a local disk or is not configured to use its local disk. Diskless nodes boot through the network, using the master node as a boot server. Diskless nodes run customer applications locally, but store and access reliable (highly available) data through the cluster network on a disk partition of the master node, using the Reliable File Service of the Foundation Services. A diskless node cannot be the master node or the vice-master node. For a listing of all other node types, see also, node.</p>
disqualified	<p>A master-eligible node that cannot participate in an election for the master role or vice-master role. Disqualified is a state used by the Cluster Membership Manager.</p>

Distributed Replicated Block Device	DRBD is open source software for data replication over IP. This software is used on Linux by the Reliable File Service when data is shared between master nodes using IP replication. On the Solaris OS, SNDR is similar to DRBD. See also, IP replication, Reliable File Service, and Sun StorEdge™ Network Data Replicator.
distributed services	Services that run on all peer nodes. The distributed services include the Cluster Membership Manager, the Node Management Agent, and the Process Monitor Daemon. See also, highly available services.
domainid	(domain identification number) The number that identifies a cluster. The <i>domainid</i> is the same for each peer node in a cluster. Also called clusterid. See also, nodeid.
double fault	The simultaneous failure of both master-eligible nodes, or the simultaneous failure of both redundant networks.
downtime	The percentage of time that a system is unavailable, including attributable, scheduled, unscheduled, total, and partial system outages. Downtime also includes outages caused by operational error. See also, availability and outage.
DRBD	Distributed Replicated Block Device
Dynamic Host Configuration Protocol	DHCP is a TCP/IP protocol that enables machines to obtain IP addresses statically or dynamically from centrally administered servers.

E

EAM	External Address Manager
eligible	See master-eligible node.
Ethernet address	A 48-bit address used to direct datalink layer transactions. The Ethernet address is synonymous with the MAC address.
external address	An IP address visible on an external network, that is assigned to a logical interface or physical interface on a peer node. See also, floating external address.
external network	A network that is not the intra-cluster network and from which cluster nodes can be reached. An external network can be Ethernet, ATM, or any other supported network type. See also, cluster network.

External Address

Manager EAM is a service that configures floating external addresses for external interfaces on the master node. After a failover or switchover, the EAM configures floating external addresses for external interfaces on the new master node. EAM also monitors the state of either the IPMP groups on the Solaris OS or the bonding interfaces on the Linux OS, and triggers a failover when an IPMP group or a bonding interface fails. See also, bonding, bonding interface, floating external address, Internet Protocol Multipathing, master-eligible node, failover, and switchover.

F

failover The transfer of services from the master node to the vice-master node when the master node fails. The vice-master node must have all of the necessary state information to take over at the moment of failover. The vice-master node expects no cooperation or coordination from the failed master node. See also, switchover.

FC-AL Fibre Channel-Arbitrated Loop

floating address triplet A logical address triplet for the node holding the master role. Diskless nodes and dataless nodes access services and data on the master node through the floating address triplet. See also address triplet.

floating external address A logical IP address visible on an external network that is assigned to an interface on the node holding the master role.

Foundation Services The basic set of services provided by the Netra HA Suite software, which includes the following Foundation Services:

- Reliable Network Communication Service (CGTP)
- External Address Manager (EAM)
- Reliable File Service (RNFS)
- Cluster Membership Manager (CMM)
- Reliable Boot Service (RBS)
- Process Monitor Daemon (PMD)
- Node State Manager (NSM)
- Node Management Agent (NMA)

Some services are distributed services, others are highly available services. See also, distributed services and highly available services.

G

guest domain Uses services from the I/O and service domains and is managed by the control domain.

H

HA-aware applications Applications that are aware of the availability of nodes and resources. The resources can be other application components, communication end points, processes, or groups of resources. An application that is HA-aware is designed to recover from a failure.

HA-unaware applications Applications that are unaware of the availability of nodes and resources. To enable these applications to recover from a failure, a dedicated agent (interfacing with CMM) must be provided to manage application recovery in the event a failure occurs. NSM can be used to write such an agent.

heartbeat An IP packet that is periodically multicast over the cluster network through each of the two physical interfaces of each peer node. When a heartbeat is detected through a physical interface, it indicates that the node is reachable and that the physical interface is alive. A probe mechanism is another way to design this mechanism.

highly available services Services that run on the master node only. The Reliable Boot Service and Reliable File Service are highly available services. If the master node or one of these services on the master node fails, a failover occurs. See also, distributed services.

horizontal scalability The ability to add client (master-ineligible) nodes to a cluster to increase the processing capacity of the system.

host part The second part of an Internet address. The host part of an IP address identifies a node on a given network. See also netmask and network part.

hot-swap The removal and replacement of a hardware component or board without shutting down the entire system.

HPI Hardware Platform Interface from Service Availability Forum (SA Forum)

hypervisor firmware layer interposed between the operating system and the hardware layer

I

- in node** A peer node that is using the Foundation Services and can communicate with other peer nodes. An “in node” can be a master-eligible node, a diskless node, or a dataless node. The master node and vice-master node must always be “in nodes.” See also, out node. For a listing of all other node types, see also, node.
- installation server** The hardware required to install the OS, the Netra HA Suite software, and some known patches on a Foundation Services cluster. The installation server can be any system capable of running the Solaris or Linux Operating Systems. Refer to the *Netra High Availability Suite 3.0 1/08 Foundation Services Getting Started Guide* for information about which OS combinations are supported between the installation server and the cluster.
- integrity** In this context, one of the proposed data management policies. See also, data management policy.
- Internet Protocol Multipathing** On the Solaris OS, IPMP is a service that allows the failover of several addresses from a failed interface to a working one. Addresses move between interfaces belonging to the same IPMP group. The floating external address can be managed by IPMP. See also, bonding, bonding interface, External Address Manager, and floating external address.
- IP** Internet Protocol
- IP replication** The copying of data from the master node to the vice-master node through the cluster network (sending data over the IP protocol). Through IP replication, the vice-master node keeps an up-to-date copy of the data on the master node. IP replication is also known as shared nothing mechanism, as it is a way to keep master and vice-master synchronized without sharing any hardware (disks). See also, replicated partition and synchronization.
- IPMP** IP multipathing

J

Java Dynamic Management Kit

JDMK is the fundamental framework for developing distributed Java™ management applications.

Java Management Extensions

JMX are tools for building distributed, Web-based, modular and dynamic solutions for managing and monitoring devices, applications, and service-driven networks.

L

ldm(1M) Logical Domain Manager utility

ldmd Logical Domains Manager daemon

local partition A partition located on a disk that is private to the node (also known as a local disk), such as a boot partition located on the system disk. For a listing of all other partition types, see also partition.

logical address A supplementary address on a physical interface or virtual interface of a node. One physical interface or virtual interface can have many logical addresses. See also, physical address and virtual address.

logical domain discrete logical grouping with its own operating system, resources, and identity within a single computer system. Also referred to as *LDom*.

Logical Domains (LDoms) Manager provides a CLI to create and manage logical domains and allocate resources to domains

logical interface A logical interface configured on a physical interface or virtual interface of a node. A logical interface on an `hme0` or `cgtp0` interface might be called `hme0:1` or `cgtp0:1`.

Logical Volume Manager LVM is an open source software for disk volume management on Linux. This is a software module for Reliable File Service when using shared disk(s). On the Solaris OS, SVM is similar to LVM.

LVM Logical Volume Manager

M

MAC address Medium Access Control address

master node A peer node (master eligible) that coordinates all cluster membership. The primary instance of External Address Manager, Reliable File Service, Reliable Boot Service, and Node State Manager run on the master node. If a floating

	external address is to be used, it must be defined on this node. There is only one master node in a cluster. For a listing of all other node types, see also, node.
master-eligible node	A peer node that can be elected as the master node or the vice-master node. Master-eligible nodes are always diskfull. For a listing of all other node types, see also, node.
master-ineligible node	A peer node that <i>cannot</i> perform the role of master node or vice-master node. Also called the client node or satellite node. Master-ineligible nodes are always diskless or dataless. For a listing of all other node types, see also, node.
MBean	(management bean) A Java interface that conforms to design patterns that expose attributes and operations to the Node Management Agent (NMA). These attributes and operations enable the NMA to recognize and manage the MBean. MBeans are defined by the JMX™ specification. For more information, go to http://java.sun.com .
MEN	Master Eligible Node
metadevice	A volume manager virtual device, referred to as a logical volume on Linux. There are several types of metadevices. Netra HA Suite software makes use of two of them: <ul style="list-style-type: none"> ■ Soft or logical metadevice: a virtual division of one or more disks (also called a soft partition) ■ Mirror metadevice: a metadevice grouping two or more sub-mirror metadevices located on different disks and mirroring the same data See also, mirror partition, slice, and soft partition.
MIB	Management Information Base
mirror partition	A partition mirror of a local, replicated, or shared partition. On the Solaris OS, you can only define a mirror partition using SVM. On Linux, you can only define a mirror partition using Linux software RAID. For a listing of all other partition types, see also, partition.
MontaVista Carrier Grade Edition Linux	MV CGE Linux is one of three versions of the Linux distribution provided by MontaVista. See http://www.mvista.com for more information.
MV CGE	MontaVista Carrier Grade Edition

N

nametag Corresponds to a daemon or a group of daemons launched by a startup script. See also, probe mechanism and startup script.

netmask A number that defines how many of the upper bits of an IP address identify a network. In class C addresses, the netmask is 24 bits. The upper 24 bits of a class C address identify the network, and the lower 8 bits identify the node. See also, network part and host part.

network See cluster network and external network.

network part The upper part of an IP address. The network part of an IP address identifies the network on which a node resides. See also, host part and netmask.

NFS Network File System

nhinstall A tool that facilitates the installation and configuration of the Netra HA Suite software and, optionally, the OS and patches on a cluster. See also, installation server.

NIC Network Interface Card

NMA Node Management Agent

NMEN Non-master-eligible node (also called master-ineligible node)

node The basic unit of hardware in a cluster on which the Foundation Services run. A node is either a blade from a blade server or a rack-mounted server. A node can also contain some secondary storage in the form of flash or compact flash memory.

See also the following types of nodes: client node, dataless node, diskfull node, diskless node, in node, master node, master-eligible node, master-ineligible node, out node, peer node, nonpeer node, predefined node, satellite node, server node, and vice-master node.

node address triplet See address triplet.

nodeid (node identification number) The number that identifies a node within a cluster. See also, clusterid and domainid.

Node Management Agent

NMA is a service that monitors cluster statistics, initiates a switchover, changes the recovery response for daemon failure, and listens for notifications of some cluster events. The NMA is a JMX compliant management agent based on the Java Dynamic Management Kit.

Node State Manager NSM is a service that executes user-provided scripts when it becomes master or vicemaster. See also, external address, master node, failover, and switchover.

nonpeer node	A node that is not configured as a member of a cluster. A nonpeer node can communicate with one or more peer nodes to access resources or services provided by a cluster. See also, peer node. For a listing of all other node types, see also, node.
nonshared packages	Packages that are not shared between the master node and the vice-master node. Nonshared packages must be added to both master-eligible nodes. See also, shared packages.
notification	An information message sent by the <code>nhcmmmd</code> daemon on a local node to services or applications registered to receive it.
NSM	Node State Manager

O

outage	An event that impairs the ability of a system to operate at its rated capacity for more than 30 seconds.
out node	A peer node on which the Foundation Services are not running and that cannot be reached by the master node. See also in node. For a listing of all other node types, see also, node.

P

partition	A physical division of a disk or a virtual division of one or more disks. For virtual divisions of disks, partitions are always referred to as soft partitions, which are used through a volume manager. See also, bitmap partition, data partition, local partition, mirror partition, replicated partition, shared partition, and soft partition.
peer node	A node in a cluster that is configured to run the Foundation Services. See also nonpeer node. For a listing of all other node types, see also, node.
physical address	The IP address of a physical interface. See also, physical interface.
physical interface	A network interface card (NIC). The name of the interface depends on the type of network card that is used. A node could have network interfaces called <code>hme0</code> and <code>hme1</code> .
PMC	PCI Mezzanine Card
PMD	Process Monitor Daemon

PMS	Platform Management Services.
predefined node	A node that is defined as part of a cluster but that has not been physically connected to the cluster. A predefined node is a peer node with an out-of-cluster role. For a listing of all other node types, see also, node.
Process Monitor Daemon	PMD is a service that monitors daemons launched by startup scripts. These daemons can be from Netra HA Suite, operating systems, or companion products. See also, startup script.
probe mechanism	See, heartbeat.

R

RBS	Reliable Boot Service
recovery	The restoration of system operation to full capacity or partial capacity. During recovery, a system performs the following tasks: <ul style="list-style-type: none"> ■ Reassigns resources to restore the system to operational state ■ Repairs any parts that were adversely affected by the failure
RedHat Package Manager	RPM Package Manager is a package management system primarily intended for Linux. RPM refers both to a package format, and to a package installation tool.
redundancy	The provision of a backup node to take over in the event of failure. The Foundation Services use the 2N redundancy model. See also, availability, reliability, and serviceability.
relative priority	The priority of one of the Foundation Services processes compared to another. In descending order of priority, Netra HA Suite daemons have the following relative priority: nhpmd > nhwdtd > nhprobed > nhcmmd > nhcrfsd > nheamd > nhnsmd. See also, actual priority and base priority.
reliability	The measure of continuous system uptime. See also availability, redundancy, and serviceability.
Reliable Boot Service	RBS is a service that uses the Dynamic Host Configuration Protocol (DHCP) and other Netra HA Suite services to ensure the boot of diskless nodes regardless of software failures or hardware failures.
Reliable File Service	RFS is a service that provides a mounted file system to make data on the master node accessible to other cluster nodes. Reliable File Service replicates (with IP replication) or mirrors (with shared disks) disk-based data on the

	master node to the vice-master node and reconfigures the floating address triplet after failover or switchover. See also, Distributed Replicated Block Device, IP replication, shared disk, and Sun StorEdge™ Network Data Replicator.
replicated partition	A partition located on local disk(s) of a master-eligible node (MEN) and replicated on local disk(s) of the other MEN. The remote copy is kept synchronized by replicating data through IP over the cluster network. See also, IP replication. For a listing of all other partition types, see also, partition.
replication	See IP replication.
RNFS	Reliable NFS
role	A membership role allocated by the Cluster Membership Manager. See also, in node, out node, master node, vice-master node.
rolling upgrade	Feature that allows an upgrade from one version of the Netra HA Suite software to a newer version, one cluster at a time, without needing to take any cluster offline. Note: This feature is not supported in versions previous to 2.1 6/03, and cannot be used to upgrade the operating system.
RPM	Red Hat Package Manager

S

SAF	Service Availability™ Forum
satellite node	See master-ineligible node.
scheduling parameter	A parameter that represents the priority with which competing services run. The scheduling parameter is defined in the <code>nhfs.conf</code> file. See also actual priority, base priority, and relative priority.
scoreboard bitmap	Bitmap in memory or bitmap in a partition (bitmap partition) that identifies which data blocks on a replicated partition are modified while the vice-master node is out of service. See also IP replication and bitmap partition. Only applicable on the Solaris OS when using the SNDR tool. See also, Sun StorEdge™ Network Data Replicator.
SCSI	Small Computer System Interface
server node	See master-eligible node.

Service Availability

Forum The SA Forum is a consortium of communications and computing companies working together to develop and publish high availability and management software interface specifications. For more information, see <http://www.saforum.org>

service domain logical domain that provides devices such as virtual switches, virtual console connectors, and virtual disk servers to other logical domains.

serviceability The probability that a service can be restored within a specified period of time following a service failure. See also availability, redundancy, and reliability.

shared disk Disk or disks in a multi-hosted disk bay separate from the master and vice-master nodes, containing data shared by the cluster. This is an alternative to IP replication for sharing data between master and vice-master nodes. See also, IP replication and Reliable File Service.

shared packages Packages that are located on the master node and replicated on the vice-master node, or located on the shared disk. Shared packages can include the Netra HA Suite packages and user applications. See also, nonshared packages.

shared partition A soft partition located on a shared disk (a disk accessible from both master-eligible nodes alternatively, also known as a dual-hosted disk). You can only define a shared partition using a volume manager, for example, Solaris Volume Manager (SVM) on the Solaris OS or Logical Volume Manager (LVM) on Linux. For a listing of all other partition types, see also, partition.

single fault The failure of a master-eligible node, the failure of a service, or the failure of one of the redundant networks. See also, double fault.

slice If a volume manager is not used, the terms slice and partition refer interchangeably to a physical division of a disk. If a volume manager is used, a slice is equivalent to a metadvice that allows access to a file system. In this case, it can refer either to soft or mirror partitions. See also, partition, soft partition, mirror partition, and metadvice.

SMS System Management Services from Service Availability Forum (SA Forum)

SNDR Sun StorEdge™ Network Data Replicator

SNMP Simple Network Management Protocol

soft partition A virtual division of one or more disks. You can only define a soft partition using a volume manager, for example, SVM on the Solaris OS or LVM on a Linux distribution. See also, metadvice and slice. For a listing of all other partition types, see also, partition.

split brain An error scenario in which a cluster has two master nodes.

SS10 Sun Studio 10 development tools

stale cluster An error scenario in which a peer node does not receive information from the master node for more than 10 seconds.

standalone CGTP	CGTP running on nonpeer nodes. See also, Carrier Grade Transport Protocol and nonpeer node.
startup script	A script that launches the Netra HA Suite daemons, some operating system daemons, and some companion product daemons. See also, probe mechanism and nametag.
subnet	See network part.
Sun StorEdge™ Network Data	
Replicator	Part of the AVS software suite, SNDR is used to provide IP replication on the Solaris OS. See also, Distributed Replicated Block Device, IP replication, Reliable File Service.
SVM	Solaris Volume Manager
switchover	The planned transfer of a service from the master node to the vice-master node. A switchover can be initiated for the purpose of system recovery or system administration. A switchover is not linked to node failure. See also, failover.
synchronization	The copying of data from the master node to the vice-master node after the vice-master node has been out of service. Synchronization occurs after startup, disk change, failover, or switchover, unless you have delayed the start of synchronization by setting the <i>RNFS_EnableSync</i> parameter to <i>FALSE</i> in the <i>nhfs.conf</i> file. See also, IP replication.
synchronized	A state where the replicated partitions on the disks of the master node and vice-master node contain exactly the same data.

T

TCP	Transmission Control Protocol
terminal server	A console access device that connects the console ports of several nodes to a TCP/IP network.

U

UDP	User Datagram Protocol
------------	------------------------

unsynchronized After a disk change, cluster startup, failover, or switchover, the master node and vice-master node disks are unsynchronized. That is, the replicated partitions do not contain the same data. See also, synchronization.

V

vertical scalability The ability to add processors and memory to individual nodes to increase their capability. See also horizontal scalability.

vice-master node A peer node (master-eligible) that has a copy of all the information on the master node. The vice-master node is available to become the master node in the event of a failover or switchover. See also, master node, master-eligible node, master-ineligible node. For a listing of all other node types, see also, node.

virtual address The IP address of the CGTP interface. See also, virtual physical interface.

Virtual Local Area Network

VLANs are software defined groups of hosts on a local area network (LAN) that communicate as if they were on the same wire, even though they are physically on different LAN segments throughout a site.

virtual logical interface

A logical interface associated with a virtual physical interface such as the CGTP interface. A logical interface associated to the CGTP interface can be called `cgtp0:1`. One virtual physical interface can have many virtual logical interfaces. See also, virtual physical interface and CGTP interface.

virtual physical interface

The CGTP interface. The CGTP interface of a node can be called `cgtp0`. See also virtual logical interface.

VLAN Virtual LAN

W

WDT Watch Dog Timer

**Wind River Platform
for Network
Equipment, Linux**

Edition One of the Linux distributions provided by Wind River, also called Wind River CGL (Carrier Grade Linux). For more information, see http://www.windriver.com/products/platforms/network_equipment/

WR CGL Wind River Carrier Grade Linux

