



# Netra™ High Availability Suite 3.0 1/08 Foundation Services Getting Started Guide

---

Sun Microsystems, Inc.  
[www.sun.com](http://www.sun.com)

Part No. 819-5241-13  
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, Sun StorageTek, Java, docs.sun.com, Netra, Solaris JumpStart, 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, Sun StorageTek, Java, docs.sun.com, Netra, Solaris JumpStart, 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**

- 1. Planning the Netra HA Suite Foundation Services Environment 1**
  - Hardware and Software Requirements for a Netra HA Suite Cluster 2
  - Using Logical Domains 3
  - Hardware and Software Requirements for the Installation Server 4
  - Hardware and Software Requirements for the Development Host 4
  - Choosing an Installation Method 5
  - Installation Task Map 6
- 2. Choosing Hardware and Software for the Cluster and the Installation Server 9**
  - Choosing Hardware for a Netra HA Suite Cluster 9
    - Choosing Hardware for Evaluation Purposes 10
      - Two-Node Cluster 10
      - Four-Node Cluster 11
    - Choosing Hardware for Real Application Testing 13
    - Sample Memory Usage for an 18-Node Cluster 14
  - Choosing Software (OS) for a Netra HA Suite Cluster 15
  - Choosing Hardware and Software for the Installation Server 16
  - Choosing a Development Host 17



# Preface

---

The *Netra High Availability Suite 3.0 1/08 Foundation Services Getting Started Guide* provides the basic information to set up the hardware and software required to install and run the Netra™ High Availability (HA) Suite 3.0 1/08 Foundation Services.

This book outlines the tasks to be accomplished to successfully install the Netra HA Suite 3.0 Foundation Services, but does not describe all of the installation and configuration options that are available with the Foundation Services.

---

## Who Should Use This Book

This book is for new Netra HA Suite customers (or evaluators) who must understand which hardware and software they need to install to use the Netra HA Suite Foundation Services and the tasks they have to go through before starting the installation of the product.

---

## Before You Read This Book

It is recommended that you read the *Netra High Availability Suite 3.0 1/08 Foundation Services Overview* for a general understanding of the product and the services provided, including highly available and distributed services.

For definitions of the terms used in this guide, refer to the *Netra High Availability Suite 3.0 1/08 Foundation Services Glossary*.

---

# How This Book Is Organized

[Chapter 1](#) provides a brief overview of the hardware and software required for the Netra HA Suite product. This chapter describes the installation methods available for installing the Netra HA Suite software, and it maps the tasks involved in installing the cluster and installation hardware.

[Chapter 2](#) provides example hardware configurations in typical clusters.

---

# Using UNIX Commands

This document might not contain information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices. Refer to the following for this information:

- Software documentation that you received with your system
- Solaris™ Operating System documentation, which is at

<http://docs.sun.com>

---

# Shell Prompts

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

---

# Typographic Conventions

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

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

---

## Related Documentation

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

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

Application	Title	Part Number
Late-breaking news	<i>Netra High Availability Suite 3.0 1/08 Release Notes</i>	819-5249-14
Introduction to concepts	<i>Netra High Availability Suite 3.0 1/08 Foundation Services Overview</i>	819-5240-13
Basic setup, supported hardware, and configurations	<i>Netra High Availability Suite 3.0 1/08 Foundation Services Getting Started Guide</i>	819-5241-13
Automated installation methods	<i>Netra High Availability Suite 3.0 1/08 Foundation Services Installation Guide</i>	819-5242-13
Detailed installation methods	<i>Netra High Availability Suite 3.0 1/08 Foundation Services Manual Installation Guide for the Solaris OS</i>	819-5237-13
Cluster administration	<i>Netra High Availability Suite 3.0 1/08 Foundation Services Cluster Administration Guide</i>	819-5235-13

<b>Application</b>	<b>Title</b>	<b>Part Number</b>
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 Getting Started Guide*, part number 819-5241-13.



# Planning the Netra HA Suite Foundation Services Environment

---

The Netra High Availability (HA) Suite Foundation Services enable you to create applications with highly available and distributed services in a loosely coupled cluster environment.

Planning the hardware, software, and networking components for a cluster configuration is an essential first step toward understanding the environment that is to be set up. This planning helps ensure a smooth installation and utilization of your cluster.

This chapter provides a high-level view of the hardware and software that is required for installing, running, and using the Netra HA Suite Foundation Services for application development.

For information about hardware and software requirements, and installation methods, see the following sections:

- [“Hardware and Software Requirements for a Netra HA Suite Cluster” on page 2](#)
- [“Using Logical Domains” on page 3](#)
- [“Hardware and Software Requirements for the Installation Server” on page 4](#)
- [“Hardware and Software Requirements for the Development Host” on page 4](#)
- [“Choosing an Installation Method” on page 5](#)
- [“Installation Task Map” on page 6](#)

---

# Hardware and Software Requirements for a Netra HA Suite Cluster

The hardware and software you choose to run the Netra HA Suite Foundation Services is directly influenced by the objectives of your installation. For example, if you are only evaluating the Netra HA Suite Foundation Services, the hardware and software requirements might be very different than they are if you are developing and testing an application using the Netra HA Suite Foundation Services.

In the latter case, your hardware and software choices will be influenced by the type of application you want to develop and deploy, your expectations for the system's performance, and the environment of your future deployments. For example, consider whether your application will be deployed on a network equipment provider [NEP] core network, in a datacenter, or in difficult field environments.

When selecting hardware and software for use on a cluster that is running the Netra HA Suite Foundation Services consider the following questions:

- Are you installing a cluster for evaluation purposes or are you developing and testing an application using the Netra HA Suite Foundation Services?
- Are you running a client/server application? If so, must your clients be part of the cluster?
- Are multiple instances of the same application servicing external devices?
- Does your application perform a high number of disk I/Os or is there a lot of network traffic inside or outside of the cluster?
- Does your application access only local data, or does it also access data that must be highly available?
- Is part of your application already available on the Solaris OS or under a Linux distribution?
- Do you want to consolidate multiple rack mount servers in a single ATCA blade server?
- Do you plan to use virtualization & share resources between logical domains on a single server?

Answering the preceding questions and likely several others will guide you as you define the cluster you will set up for your application and on which you will use the Netra HA Suite Foundation Services.

Depending on your answers to the preceding questions, you will decide whether or not to set up a cluster with client nodes (diskless or dataless), whether to support data replication over IP or through shared disks, whether or not to handle external access lines, and whether to run the Solaris OS or a Linux distribution. You will also

decide if you want to install your cluster on rackmounted servers or blade servers, and whether or not your solution needs to be NEPS-compliant. If you plan to use virtualization, it implies a choice of well-identified hardware (based on the UltraSPARC® T1 processor family) and limits the operating system choice.

At a minimum, a cluster running the Netra HA Suite 3.0 software requires the following:

- Two server nodes (master-eligible nodes) with their local disks
- Two network interfaces on each node
- Two Ethernet switches
- One terminal server

Other hardware components are optional and should be selected based on the answers you provided to the previously listed questions. Some of these optional components include the following:

- Client nodes (diskless or dataless), up to 62
- SCSI disk bay or network attached storage (NAS) if using shared disk for data replication
- Supplementary network interface cards if external access is required

[Chapter 2](#) provides examples of hardware configurations in typical clusters. Along with each of this typical cluster, some hints are given about the type of Application such a cluster will preferably support.

Running the Netra HA Suite 3.0 Foundation Services on your cluster requires either the Solaris Operating System or a Carrier Grade Linux distribution. But if you choose to run Netra HA Suite Foundation Services on a cluster where virtualization and partitioning technology are in use, only the Solaris OS can be installed. For information about supported Solaris releases and Linux distributions, see *Netra High Availability Suite 3.0 1/08 Release Notes*.

---

## Using Logical Domains

Sun's Logical Domains (LDom)s technology is a server virtualization and partitioning technology that enables the allocation of various system resources, such as memory, CPUs, I/O, and storage into partitions known as logical or virtual domains. Each logical domain can have an independent operating system, resources, and identity within a single computer system. Specialized service and control domains allow these resources to be managed using the Logical Domains Manager software.

For information about the LDom configurations that are supported with this release of the Netra HA Suite Foundation Services, see the *Netra High Availability Suite 3.0 1/08 Release Notes*.

---

**Note** – For more information about LDom, refer to the *Logical Domains (LDoms) Administration Guide* (819-6428) and the Sun BluePrints™ Online article, “Beginners Guide To LDom: Understanding and Deploying Logical Domains” (820-0832).

---

---

## Hardware and Software Requirements for the Installation Server

To install the OS and Netra HA Suite 3.0 software on your cluster, you need an installation server and a network link between the installation server and your cluster. The installation server can be any hardware that is capable of running the Solaris OS or a Linux distribution.

See the *Netra High Availability Suite 3.0 1/08 Foundation Services Installation Guide* for information about choosing the hardware and software for the installation server, connecting the cluster and installation server hardware, and installing and configuring the cluster software from the installation server.

---

## Hardware and Software Requirements for the Development Host

If you are developing applications that you plan to deploy on a cluster running the Netra HA Suite 3.0 Foundation Services, you can choose to install a development host.

The hardware required to support a development host is optional, but can be any hardware that is capable of running the Solaris OS or a Linux distribution.

See the *Netra High Availability Suite 3.0 1/08 Foundation Services Installation Guide* for information about installing a development host.

---

**Note** – The installation server can also be used as a development host, especially when you are only performing an evaluation of the Netra HA Suite 3.0 product.

---

---

# Choosing an Installation Method

The Netra HA Suite product provides two ways of installing software on the cluster.

- *Automated installation with the `nhinstall` tool.* The `nhinstall` tool, running on the installation server, enables you to install the OS of your choice and the Netra HA Suite software on a cluster, regardless of the hardware configuration of the cluster. It also automatically manages the configuration of the cluster, based on predefined configuration files.

This tool is flexible and provides various configuration options that you can adapt to your requirements. For more information, see the *Netra High Availability Suite 3.0 1/08 Foundation Services Installation Guide*.

- *Manual installation.* You can manually install the software on a cluster, regardless of the hardware configuration in use. A manual installation requires you to manage, step-by-step, the installation of the software (the OS and the Netra HA Suite software) on the cluster, without any help from an automated tool.

Manual installation provides greater flexibility when installing various components of the Foundation Services. However, manual installation can result in a cluster configuration that is not easily reproducible on other clusters.

Manual installation is not recommended for use on the Solaris 10 OS and Linux because the procedures involved are complex and can be a source of errors. For more information, see the *Netra High Availability Suite 3.0 1/08 Foundation Services Manual Installation Guide for the Solaris OS*.

If you install Netra HA Suite Foundation Services on a cluster where the virtualization and partitioning technology are in use, Solaris OS and LDOMs software have to be installed manually first (see the *Netra High Availability Suite 3.0 1/08 Foundation Services Installation Guide* for detailed instructions). Then, you can choose an automated installation or a manual installation to install Netra HA Suite on the cluster.

# Installation Task Map

The following table outlines the tasks for choosing and installing the cluster hardware and the installation hardware.

Task	Description	For Instructions
Choose the hardware configuration for the cluster that is going to run the Netra HA Suite Foundation Services and define the network topology.	Choose the number and type of nodes of the cluster. Try to include nodes that you might want to add to the cluster in the future.	<a href="#">“Choosing Hardware for a Netra HA Suite Cluster” on page 9</a>
Choose the software (operating system [OS]) that will run on the cluster nodes.	Choose the OS and supplementary software to install on nodes of the cluster.	<a href="#">“Choosing Software (OS) for a Netra HA Suite Cluster” on page 15</a>
Choose the hardware for the installation server.	Choose the hardware to be used for the installation server and decide how to connect it to the cluster.	<a href="#">“Choosing Hardware and Software for the Installation Server” on page 16</a>
Choose the software for the installation server.	Choose the software to be installed on the installation server	<a href="#">“Choosing Hardware and Software for the Installation Server” on page 16</a>
Install the cluster hardware and its network topology.	Install the cluster nodes, the Ethernet switches, and the terminal server.	Hardware documentation and the <i>Netra High Availability Suite 3.0 1/08 Foundation Services Installation Guide</i>
Install the installation server hardware.	The installation server is required for any installation method you choose.	Hardware documentation and the <i>Netra High Availability Suite 3.0 1/08 Foundation Services Installation Guide</i>
Connect the installation server hardware and the cluster hardware.		<i>Netra High Availability Suite 3.0 1/08 Foundation Services Installation Guide</i>
Install the installation server software.	The installation server is required for any installation method you choose.	<i>Netra High Availability Suite 3.0 1/08 Foundation Services Installation Guide</i>
Install the OS and Netra HA Suite Foundation Services software on the cluster nodes.	Use the nhinstall tool if you choose to perform an automated installation.	<i>Netra High Availability Suite 3.0 1/08 Foundation Services Installation Guide</i>



If you are planning to develop applications using the Netra HA Suite software API, you must also perform the following tasks to install a development server:

Task	Description	For Instructions
(optional) Choose the hardware for the development server and decide how to connect it to the cluster.	Perform this task only if you are planning to develop applications using the Netra HA Suite software API.	<a href="#">“Choosing a Development Host” on page 17</a>
(optional) Choose the software for the development server.	If you are installing a development server, decide which software to install on it.	<a href="#">“Choosing a Development Host” on page 17</a>
(optional) Install the development server HW and connect it to the cluster.	If you are installing a development server, connect the cluster nodes to the Ethernet switches and the terminal server. Connect the installation server.	Hardware documentation and the <i>Netra High Availability Suite 3.0 1/08 Foundation Services Installation Guide</i>
(optional) Install the development server software.	Connect the cluster nodes to the Ethernet switches and the terminal server. Connect the installation server.	<i>Netra High Availability Suite 3.0 1/08 Foundation Services Installation Guide</i>



## Choosing Hardware and Software for the Cluster and the Installation Server

---

Plan the installation and configuration of your cluster thoroughly to avoid setbacks and delays. Choose the size and hardware configuration of your cluster to suit your purpose. The following sections describe the considerations involved in making these choices:

- [“Choosing Hardware for a Netra HA Suite Cluster” on page 9](#)
- [“Choosing Software \(OS\) for a Netra HA Suite Cluster” on page 15](#)
- [“Choosing Hardware and Software for the Installation Server” on page 16](#)

---

## Choosing Hardware for a Netra HA Suite Cluster

Configurations provided in this section are examples of supported configurations. If you plan to install a hardware configuration that is not described in the following examples, especially if you are using a mix of hardware, ask your support team for information about the supported configurations and available configuration options.

# Choosing Hardware for Evaluation Purposes

Perform evaluations of the Netra HA Suite software on a cluster that is easy to set up, for example, either a two-node or four-node cluster. Use a two-node cluster (master-eligible nodes only) if you do not plan to use client (diskless or dataless) nodes in your cluster with your application. Use a four-node cluster (two master-eligible nodes and two master-ineligible nodes) in other cases.

## Two-Node Cluster

The quickest setup for a two-node cluster is to use rackmounted servers. Whichever rackmounted server you use, the cluster must be configured as follows (with IP replication used for data sharing between the two master-eligible nodes):

- Two rackmounted servers configured as master-eligible nodes.
- Some partitions of the rackmounted server's internal disks configured as replicated partitions (for storage of highly available data).
- On-board Gigabit Ethernet interfaces (or supplementary cards) of the two rack-mounted servers for the configuration of the cluster network and (optionally) for providing external access to the cluster. In the latter case, use at least four NICs.
- Two external Ethernet switches (see Chapter 1 of the *Netra High Availability Suite 3.0 1/08 Foundation Services Installation Guide* for information about configuring and connecting the switches).
- One terminal server to manage the consoles (see Chapter 1 of the *Netra High Availability Suite 3.0 1/08 Foundation Services Installation Guide* for information about configuring and connecting the terminal server).

The following list provides some examples of a two-node cluster based on rack-mounted servers:

- Two-node cluster with two Netra 120 or two SunFire V210 servers
- Two-node cluster with two Netra 240 or two SunFire V240 servers
- Two-node cluster with two Netra 440 or two SunFire V440 servers
- Two-node cluster with two Netra T2000 servers

Choose Netra T2000 servers if you want to test the behavior and performance of the Netra HA Suite Foundation Services on CMT-based hardware.

---

**Note** – If you need to evaluate data sharing between two MENs by using shared disks (instead of IP replication), add an external disk array to your configuration. Check with your support team to determine which external disks are supported for use on a particular server.

---

---

**Note** – All of the servers listed in this section are running with only the Solaris OS. To evaluate the Foundations Services on Linux, you must set up a cluster with two Netra CP3020 blades (Opteron-based blades) in an ATCA chassis (preferably the Sun Netra CT900 chassis).

---

## Four-Node Cluster

If you already have a two-node cluster set up as described in the preceding section, and want to evaluate the testing of client nodes, you can either add new rack-mounted servers connected to the cluster network, or add an ATCA chassis with two blades to your configuration. The second option is recommended, as it offers more flexibility in the choice of which client nodes you use. For example, you can use either SPARC®, Opteron™, or CMT-based blades running the Solaris OS or Linux, and will then have the potential to test scalability afterwards (you can easily add up to 12 blades in the chassis).

A four-node cluster based on two rackmounted servers and two ATCA blades must contain the following:

- Two rackmounted servers configured as master-eligible nodes.
- One ATCA (Netra CT 900) chassis and two ATCA blades configured as master-ineligible nodes (diskless or dataless).
- Replicated partitions defined on the internal disks of the rackmounted servers (for storage of highly available data).
- On-board Gigabit Ethernet interfaces (or supplementary cards) of the rack-mounted servers for the configuration of the cluster network and (optionally) for providing external access to the cluster. In the latter case, use of at least four NICs is recommended.
- Two external Ethernet switches (see Chapter 1 of the *Netra High Availability Suite 3.0 1/08 Foundation Services Installation Guide* for information about configuring and connecting the switches).
- One terminal server to manage the consoles (see Chapter 1 of the *Netra High Availability Suite 3.0 1/08 Foundation Services Installation Guide* for information about configuring and connecting the terminal server).

The following are examples of a four-node cluster based on two rackmounted ATCA blades:

- Four-node cluster with two Netra 240 servers and one Netra CT 900 chassis containing two Netra CP3010 or two Netra CP 3020 blades.
- Four-node cluster with two Netra 440 servers and one Netra CT 900 chassis containing two Netra CP3010 or two Netra CP 3020 blades.

- Four-node cluster with two Netra T2000 servers and one Netra CT 900 chassis containing two Netra CP3060 blades.

---

**Note** – If you need to evaluate data sharing between two MENs by using shared disks (instead of IP replication), add an external disk array to your configuration. Check with your support team to determine which external disks are supported for use on a particular server.

---

---

**Note** – If you intend for your cluster to have all nodes running Linux, to evaluate the Foundations Services, you must set up a cluster with only Netra CP3020 blades (Opteron-based blades) in one or two ATCA chassis (preferably the Sun Netra CT900 chassis).

---

If you are installing a four-node cluster from scratch, use an ATCA blade server (Netra CT900 chassis with Netra CP30xx blades) to ensure maximum flexibility and possible scalability. A four-node cluster based on an ATCA blade server must be configured with the following hardware:

- One ATCA (Netra CT 900) chassis in which the four blades are in the same chassis, or two ATCA chassis in which two blades (one MEN and one NMEN) are in each chassis.
- Two ATCA blades (Netra CP30xx) configured as master-eligible nodes.
- Two ATCA blades (Netra CP30xx) configured as master-ineligible nodes (diskless or dataless).
- Replicated partitions defined on the internal disks of the two MENs (for storage of highly available data).
- Internal ATCA base or extended fabrics used for the cluster network configuration. Other Gigabit Ethernet interfaces (available on the blades) are used for connecting external networks (if required).
- Two ATCA switch blades (see Chapter 1 of the *Netra High Availability Suite 3.0 1/08 Foundation Services Installation Guide* for information about configuring and connecting the switches).
- Two ATCA shelf manager blades to manage the consoles (see Chapter 1 of the *Netra High Availability Suite 3.0 1/08 Foundation Services Installation Guide* for information about configuring and connecting the terminal server).

The following are examples of a four-node cluster based on an ATCA blade server:

- Four-node cluster with one Netra CT 900 chassis containing four Netra CP3010, four Netra CP3020, or four Netra CP3060 blades.
- Four- node cluster with one Netra CT 900 chassis containing two Netra CP3010 blades and two Netra CP3060 blades, or two Netra CP3020 blades and two Netra CP3060 blades.

- Four-node cluster with two Netra CT 900 chassis containing two Netra CP3020 or two Netra CP3060 blades in each chassis (one MEN and one NMEN in each chassis).

---

**Note** – If you plan to evaluate a four-node cluster on a system where Logical Domains are in use, build your cluster using either an ATCA blade server with two Netra CP3060 blades or two Netra rackmounted T2000 servers (both are UltraSPARC T1 processor-based hardware). On each Netra CP 3060 blade (or each Netra T2000 server) you will configure three domains (one control domain and two guest domains). Each domain will run the Solaris OS. Each guest domain will be considered as a node of the cluster. There will be one master-eligible and one master-ineligible node on each Netra CP3060 blade or Netra T2000 server.

---

## Choosing Hardware for Real Application Testing

The details provided in [“Choosing Hardware for Evaluation Purposes”](#) on page 10 also apply for testing real applications, but the size of your cluster must be aligned with your expectation in term of performance. This means that you might have to build clusters of up to 64 nodes (the maximum number of nodes currently supported by the Netra HA Suite software).

Consider the following hardware configurations based on the number of nodes you need in a cluster. These suggestions also provide an indication of memory requirements for the Foundation Services.

**TABLE 2-1** Suggested Hardware Configurations for a Netra HA Suite Cluster

Cluster Size	Recommended Hardware
2 nodes	Rackmounted servers: Netra 120/240/440/1290/T2000 and SunFire equivalent
4–12 nodes	1 or 2 ATCA chassis (Netra CT 900) up to 12 Netra CP30xx blades as MENs and NMENs Note: Netra CP3020 is mandatory if Linux is the preferred choice of OS OR 2 rackmounted servers as MENs (Netra 120/240/440/1290/T2000), 1 or 2 ATCA chassis (Netra CT 900), up to 10 Netra CP30xx blades as NMENs (diskless or dataless), Note: Netra CP3020 is mandatory if Linux is the preferred choice of OS for NMENs
12–48 nodes	2 rackmounted servers as MENs (Netra 240/440/1290/T2000), up to 4 ATCA chassis (Netra CT 900), up to 46 Netra CP30xx blades as NMENs dataless (maximum of 12 per chassis), Note: Netra CP3020 is mandatory if Linux is the preferred choice of OS for NMENs
48–64 nodes	2 rackmounted servers as MENs (Netra 440/1290/T2000), up to 6 ATCA chassis (Netra CT 900), up to 62 Netra CP30xx blades as NMENs dataless (maximum of 12 per chassis) Note: Netra CP3020 is mandatory if Linux is the preferred choice of OS for NMENs

## Sample Memory Usage for an 18-Node Cluster

Each cluster must have two master-eligible nodes. You can have a mix of diskless nodes and dataless nodes in a cluster. For definitions of the types of nodes, see the *Netra High Availability Suite 3.0 1/08 Foundation Services Glossary*.



In a sample 18-node cluster, the memory footprint of each running daemon is as follows:

**TABLE 2-2** Sample Results of Memory Usage for 18-Node Cluster

Function	Memory Used
DHCP	4.8 MB on the master and vice-master nodes
PROBE	2.6 MB on every node
CMM	4.7 MB on the master and vice-master nodes and 3.4 MB on the remaining nodes
CRFS	3.4 MB on the master and vice-master nodes
SNDR	2.7 MB on the master and vice-master nodes
JVM	50 MB (approximately)

The total memory used for Foundation Services-related daemons is approximately 70 megabytes for the master and vice-master nodes, and 55 megabytes for the remaining nodes.

# Choosing Software (OS) for a Netra HA Suite Cluster

In general, the OS that is chosen for a cluster is a strategic decision, made at the corporate level, and not a technical choice.

To get the best use of the Netra HA Suite Foundation Services, you must run them under the Solaris OS (primarily, the Solaris 10 OS), which is supported on all of the hardware referenced in this guide. Every service of the Foundation Services is available for use with it.

If you choose to use Linux, you can use only Netra CP3020 blades in an ATCA chassis to run the Netra HA Suite Foundation Services. Also, some services are not available at all under Linux (for example, diskless support). Further, some services have limitations under Linux (for example, IPv6 addresses are not supported on an external network). For information about the limitations that exist under Linux, see the *Netra High Availability Suite 3.0 1/08 Release Notes*.

If you have to use Linux to run your application, a good compromise could be to have two MENs on rackmounted servers running the Solaris 10 OS, with some NMENs (Netra CP3020s in an ATCA chassis) running your application under Linux. This configuration enables you to run your application with Linux, while benefiting from the Netra HA Suite services that are running the Solaris OS.

---

## Choosing Hardware and Software for the Installation Server

An installation server is required for all installation methods. An installation server enables you to install the operating system (Solaris or Linux) and the Netra High Availability (HA) Suite software on the cluster.

The installation server requires the following.

---

Hardware requirements	UltraSPARC® platform or i386 Sun platforms Two network devices, as follows: <ul style="list-style-type: none"><li>• If the installation server is part of the public network, one network device is used to connect the installation server to this public network. The other network device is used to connect the installation server to the cluster network.</li><li>• If the installation server is a portable machine, you require only one network device to connect to the cluster network</li></ul>
Operating system	Solaris Operating System or Linux Operating System. To install a cluster running the Solaris OS, you must install the Solaris OS on the installation server. There is no need to have the same release of Solaris on the installation server and the cluster. To install a cluster running the supported Linux distribution, you must have the Solaris 9 or 10 OS, or a SuSe 9 distribution installed on the installation server.
Software requirements	Perl Version 5, which is available with the Developer Solaris Software Group.
Disk capacity	Minimum 1.5 Gbytes for a Solaris software distribution, 4 Gbytes for an eight-node cluster.
Free space	Minimum 1.5 Gbytes after the Solaris Operating System has been installed.

---

---

# Choosing a Development Host

If you are developing applications that you plan to deploy on a cluster running the Foundation Services, you can install a *development host*. The development host is an optional hardware component. It can be on one (or more) additional servers, or the installation server can be used for the development environment, as well. If you are developing applications using the Cluster Membership Manager (CMM) API or the Service Availability Forum/Cluster Manager (SA Forum/CLM) API, you might require specific software. For more information about CMM, SA Forum/CLM, and the specific software required to develop applications for your cluster, see the *Netra High Availability Suite 3.0 1/08 Foundation Services CMM Programming Guide* and the *Netra High Availability Suite 3.0 1/08 Foundation Services SA Forum Programming Guide*.

The development host requires the following:

---

Hardware requirements	UltraSPARC platform and i386 Sun Platforms One network device
Operating system	Solaris Operating System or Linux Operating System
Software requirements	Sun™ Studio 10 software Forte™ Developer 6 Software Suite (at least Update 1) Java™ 2 Software Development Kit Standard Edition
Disk capacity	1.3—2.6 Gbytes, depending on the Solaris OS version in use
Free space	Minimum 1.5 Gbytes after the Solaris OS has been installed

---



# Index

---

## **C**

- cluster hardware
  - planning to install, 6
  - task list, 6

## **D**

- development host (optional), hardware requirements, 17

## **I**

- installation
  - manual, 5
  - method, choosing, 5
  - nhinstall tool, 5
  - planning, 5
- installation hardware
  - development host (optional), 17
  - planning, 6
  - task list, 6

## **M**

- manual installation, 5

## **N**

- nhinstall tool
  - installation, 5

