




SCO

UnixWare[®] 7 NonStop Clusters[®] Installation

Student Guide

Course Release: 1.0

For Software: UnixWare[®] 7 NonStop[®] Clusters Version 7.1.1

©2000 The Santa Cruz Operation, Inc. All rights reserved.

This publication is protected under copyright laws and international treaties.

Information in this document is subject to change without notice and does not represent a commitment on the part of The Santa Cruz Operation, Inc.

SCO documents are provided as is and may include technical inaccuracies or typographical errors. The Santa Cruz Operation, Inc. reserves the right to add, delete, change or modify SCO documents at any time without notice. SCO, the SCO logo, The Santa Cruz Operation, SCO OpenServer, Tarantella, the Tarantella logo and UnixWare are trademarks or registered trademarks of The Santa Cruz Operation, Inc. in the USA and other countries. IXI, X.desktop, Deskworks, Wintif, IXI Panorama, Deskterm, and IXI Desktop are trademarks or registered trademarks of IXI Limited, a subsidiary of The Santa Cruz Operation, Inc. Visionware, Devkit, Vision, Esprit, Kodon, Super.Vision, Visionware Super.Vision, Term.Vision, Visionware Direction, the Visionware logo, Visionware SQL-Retriever, and XVision are registered trademarks of Visionware Limited, a subsidiary of The Santa Cruz Operation, Inc. SuperVision, TermVision, X-Visionware, SQL-Retriever, and Vision Builder are trademarks of Visionware Limited. UNIX is a registered trademark of The Open Group in the United States and other countries. Cheyenne and ARCserve are registered trademarks of Cheyenne Software, Inc. Netscape, Netscape Navigator, Netscape Communications Server, Netscape Commerce Server, Netscape Proxy Server, Netscape FastTrack Server, Netscape Enterprise Server, Netscape SuiteSpot, Netscape Catalog Server, Netscape News Server, Netscape Mail Server, and Netscape Navigator Gold are trademarks of Netscape Communications Corporation. NFS was developed by Legent Corporation (formerly Lachman Associates, Inc.) based on LACHMAN SYSTEM V NFS. LACHMAN is a trademark of Legent Corporation. NFS is a trademark of Sun Microsystems, Inc. TCP/IP was developed by Legent Corporation (formerly Lachman Associates, Inc.) based on LACHMAN SYSTEM V STREAMS TCP, a joint development of Lachman Associates and Convergent Technologies. MPX was developed by Corollary, Inc. VP/ix is a product developed and licensed by Phoenix Technologies, Ltd/INTERACTIVE Systems Corporation. XRemote is a registered trademark of Network Computing Devices, Inc. Oracle is a registered trademark of Oracle Corporation, Redwood City, California. Java is a trademark of Sun Microsystems, Inc. Compaq, Tandem, Integrity, NonStop, NetFlex, ProLiant, ROMPaq, ServerNet, SmartStart, registered in United States Patent and Trademark Office. Intel and Pentium are registered trademarks, and Pentium II Xeon and Pentium III Xeon are trademarks of Intel Corporation. SCO, Veritas, VxFS, and VxVA are registered trademarks of Veritas Software Corporation. Other product names mentioned herein may be trademarks and/or registered trademarks of their respective companies. All other brand and product names are or may be trademarks of, and are used to identify products or services of, their respective owners.

The SCO software that accompanies this publication is commercial computer software and, together with any related documentation, is subject to the restrictions on US Government use as set forth below. If this procurement is for a DOD agency, the following DFAR Restricted Rights Legend applies:

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the US Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of Rights in Technical Data and Computer Software Clause at DFARS 252.227-7013. Contractor/Manufacturer is The Santa Cruz Operation, Inc., 400 Encinal Street, Santa Cruz, CA 95060.

If this procurement is for a civilian government agency, this FAR Restricted Rights Legend applies:

RESTRICTED RIGHTS LEGEND: This computer software is submitted with restricted rights under Government Contract No. _____ (and Subcontract No. _____, if appropriate). It may not be used, reproduced, or disclosed by the Government except as provided in paragraph (g)(3)(i) of FAR Clause 52.227-14 alt III or as otherwise expressly stated in the contract. Contractor/Manufacturer is The Santa Cruz Operation, Inc., 400 Encinal Street, Santa Cruz, CA 95060.

Table of Contents

Lesson 1: Course Introduction	1-1
Introduction.....	1-3
Course goal and purpose.....	1-3
Course Lessons	1-5
Target audience	1-5
Student Materials	1-6
The Student Guide	1-6
SCO documentation.....	1-6
Typographical Conventions	1-7
Lesson 2: General Tips for Building a Cluster	2-1
Tower and Rack Installations.....	2-7
Cable Management	2-8
Rack Operations.....	2-9
Rack planning	2-9
Rack assembly and loading.....	2-10
Rack cabling	2-13
Site Planning	2-14
Lesson Summary.....	2-15
Answers to Exercise and Review Questions.....	2-16
Lesson 3: Assembling and Initializing Cluster Hardware	3-1
Assembly Overview.....	3-7
Hardware Assembly.....	3-10
Assembling racks	3-11
Assembling UPSes and power strips	3-13
Assembling shared storage	3-15
Assembling nodes	3-17
Assembling SAN, LAN, and KVM switches or hubs	3-19
Assembling monitors	3-21
Hardware Initialization	3-23
Initializing nodes.....	3-23
Initializing shared storage.....	3-25
Installing the split-brain avoidance cable	3-25
Testing the SAN.....	3-26
LAN testing.....	3-26
Lesson Summary.....	3-27
Lab: Assembling Cluster Hardware.....	3-28
Cross-node mirroring procedure.....	3-29
Fibre Channel/JBOD/RSO procedure.....	3-33
Answers to Review Questions	3-37

Lesson 4: Installing SCO UnixWare 7	4-1
--	------------

Understanding Software Installation	4-6
Special HBA driver considerations.....	4-7
Licensing Information.....	4-8
UnixWare 7 licenses	4-8
NSC licenses	4-9
Other licenses.....	4-10
Cluster Virtual Host Name.....	4-11
Understanding Disk and Filesystem Configuration.....	4-12
The cluster <i>root</i> filesystem.....	4-13
Online System Upgrade (OSU)	4-14
Planning non- <i>root</i> filesystems	4-15
Vendor-specific partitions	4-15
Deferring Network Configuration	4-16
Package Selection	4-17
Using EFSes and PTFs.....	4-18
Program Temporary Fixes (PTFs).....	4-18
Lesson Summary.....	4-19
Lab: Installing UnixWare 7.1 on the Initial Root Node.....	4-20
Continuing with the installation.....	4-26
Post-installation tasks	4-26
Answers to Exercise and Review Questions.....	4-27

Lesson 5: Installing NSC Software	5-1
--	------------

Understanding NSC Software Installation.....	5-6
Preparing to Install NSC on the Initial Root Node	5-7
Applying Program Technical Fixes (PTFs).....	5-7
Encapsulating the <i>root</i> filesystem.....	5-8
Preparing shared storage.....	5-9
Removing SDI ghost devices.....	5-11
Installing NSC on the Initial Root Node.....	5-12
Selecting optional packages.....	5-12
Configuring <i>root</i> filesystem failover.....	5-13
Configuring the CVIP and PIP	5-14
Configuring optional packages	5-15
Post-Installation Tasks	5-16
Lesson Summary.....	5-17
Lab: Installing NSC on the Initial Root Node	5-18
Overview.....	5-18
Procedure for CNM clusters	5-19
Procedure for shared-storage clusters	5-27
Relevant Commands and Files	5-34
Answers to Review Questions	5-35

Lesson 6: Installing the Failover Root Node	6-1
Understanding NSC Dependent Node Installation	6-6
Installing the Failover Root Node.....	6-7
Creating dependent node boot diskettes	6-7
Booting the failover root node using the DNBDs.....	6-8
Adding the node license.....	6-8
Checking the fdisk table for CNM	6-9
Configuring external storage for dependent nodes	6-10
Configuring split-brain avoidance for CNM	6-10
Configuring networking on the dependent node.....	6-11
Installing the Mirroring license.....	6-11
Lesson Summary.....	6-12
Lab: Creating Dependent Node Boot Diskettes.....	6-13
Overview.....	6-13
Procedure: Using make_floppy to create DNBDs.....	6-14
Lab: Installing the Failover Root Node	6-15
Overview.....	6-15
Procedure: Installing the failover root node for CNM clusters	6-16
Procedure: Installing the failover root node for shared-storage clusters	6-20
Lab: Testing Failover	6-23
Overview.....	6-23
Procedure: Simulating root node failure	6-24
Relevant Commands and Files	6-26
Answers to Review Questions	6-27

1

Course Introduction



Lesson 1

Course Introduction

1.1 Introduction

This ***UnixWare 7 NonStop Clusters Installation*** course is designed for individuals who need to gain a firm understanding of how to install the hardware and software for NonStop Clusters, hereafter referred to as NSC. The target audience includes technicians that will install and configure NSC. The course may also be useful to system administrators that will maintain NSC systems on a daily basis and programmers who will develop cluster-aware applications.

1.1.1 Course goal and purpose

At the conclusion of this course, students should be able to:

- Assemble a cluster in either tower or rack form
- Perform hardware setup and diagnostics on cluster components including nodes, storage devices, and Server Area Network (SAN)
- Install UnixWare 7.1 on the initial root node in a cluster
- Install NSC on the first node in a cluster
- Install the failover root node

AV 1-1

The slide is titled "NonStop Cluster Hardware and Software Installation" in a blue-bordered box at the top. Below the title is a bulleted list of six items. At the bottom, there is a blue footer bar containing the text "Intro01" on the left and "© The Santa Cruz Operation, Inc." on the right.

- Introduction
- General Tips for Building a Cluster
- Assembling Cluster Hardware
- Installing UnixWare 7.1 on the Initial Root Node
- Installing NSC Software on the Initial Root Node
- Installing Dependent Nodes

Intro01 © The Santa Cruz Operation, Inc.

1.1.2 Course Lessons

To achieve the course goals, this course contains the following lessons:

- Introduction (this lesson) — course goals, topics, and typographical conventions
- General Tips for Building a Cluster — tips for assembling, shipping, and preparing sites for tower and rack-mount clusters
- Cluster Hardware Assembly — assembling the individual components into a cluster
- Installing UnixWare 7.1 on Node 1 — installing UnixWare 7.1 in preparation for turning the computer into node 1 of an NSC cluster
- Installing Dependent Nodes — installing NSC on the failover root node

1.1.3 Target audience

This class is aimed at people who need to become NSC installation technicians. NSC System administrators, support technicians, and developers may also find this course useful. Students are likely to have a wide variety of backgrounds in UNIX[®], UnixWare, and SCO OpenServer[™]. They must have the following prerequisite skills before attending this course:

- Passed the SCO UnixWare 7 ACE certification tests
- Passed the **UnixWare[®] 7 NonStop Clusters[®] Concepts** course test

1.2 Student Materials

This student guide is yours to keep and write notes in. Your instructor might have other reference materials in the classroom for your use while attending this class. Please do not take or write in these materials as they will be needed again for later classes.

1.2.1 The *Student Guide*

The *Student Guide* contains:

- Outline of presentation
- Note-taking space
- Copies of all transparencies

1.2.2 SCO documentation

This course refers to the following SCO printed documentation:

- *SCO UnixWare NonStop Clusters Release Notes*
- *SCO UnixWare NonStop Clusters Installation Guide*

This course also refers to the online documentation using **scohelp**(X1).

1.3 Typographical Conventions

This publication presents commands, filenames, keystrokes, and other special elements as shown here:

lp or lp(1)	Commands, device drivers, programs, and utilities (names, icons, or windows); the number in parentheses indicates the reference manual section in which the command, driver, program, or utility is documented
<i>/new/client.list</i>	Files, directories, desktops (names, icons, or windows), system, network, and printer names
root	User names
<i>arg</i> or <i>arg</i>	Placeholders (replace with appropriate name or value)
[<i>arg</i>]	Surrounds an <i>option</i> or <i>cmdarg</i> that is not mandatory
<i>arg ...</i>	Indicates multiple occurrences of the <i>option</i> or <i>cmdarg</i>
\$	Command-line prompt for commands executable from the Bourne shell by any user
#	Command-line prompt for commands only executable by the root user from the root account
F2	Keyboard keys
CTRL SPACE	Press CTRL and SPACE keys simultaneously
CTRL ALT DEL	Press all three keys, CTRL , ALT , and DEL , simultaneously
CTRL d	Press CTRL key. With CTRL still depressed, press D .
ALT 0240	Press ALT key. With ALT still depressed, press and release in sequence 0 , 2 , 4 , and 0 .
CTRL Q + i	Press CTRL key. With CTRL still depressed, press Q . Release both keys and press I .
CTRL Q + I	Press CTRL key. With CTRL still depressed, press Q . Release both keys and press SHIFT I .
Exit program?	System output (prompts, messages)

tar tv	User input, such as on a command line
“Description”	Field names or column headings (on screen or in database)
Cancel	Button names
Edit	Menu names
Copy	Menu items
File ⇨ Find ⇨ Text	Sequences of menus and menu items
open or open(2)	Library routines, system calls, kernel functions, C keywords; the number in parentheses indicates the reference manual section in which the file is documented
\$HOME	Environment or shell variables
SIGHUP	Named constants, signals, and kernel parameters
“adm3a”	Data values
<i>employees</i>	Database names and tables
buf	C program structures
b_b.errno	Structure members

2

General Tips for Building a Cluster



Lesson 2

General Tips for Building a Cluster

Introduction

Implementing clusters requires careful planning and preparation. This lesson provides tips for simplifying assembly, installation, and subsequent management and maintenance of the cluster.

Objectives

By the end of this lesson, you will be able to:

- Describe advantages of using rack or tower forms for cluster hardware
- Define rules for labeling cables for easy identification of their type, source, and destination
- Describe the guidelines for placing components within and between racks, taking into consideration power requirements, weight distribution, and heat dissipation
- Select a location for housing a cluster, including power requirements, weights, and sizes of cluster components

AV 2-1

General Tips for Building a Cluster

- Tower and Rack Installations
- Cabel Labeling
- Rack Loading
- Site Planning


Gentips01 © The Santa Cruz Operation, Inc.

Topics


In order to meet these objectives, this lesson discusses:

- Tower and Rack Installations — deciding whether to use a rack or tower form
- Cable Management — placing components to meet cabling requirements, and how to label cables
- Rack Operations — placing components within a rack, lists issues for dealing with multi-rack clusters, and lists concerns you must be aware of when working with a rack
- Site Planning — issues that must be addressed before the cluster shows up at its final destination for final assembly and installation, including how to get the cluster to that location

AV 2-2



AV 2-3



2.1 Tower and Rack Installations

As discussed in the *UnixWare® 7 NonStop Clusters® Concepts* course, an NSC cluster contains many hardware components. Depending on the required capacity, a cluster includes some or all of this equipment:

- Two complete computers, known as nodes
- Protected storage housed in one or more external cabinets
- A Keyboard/Video/Mouse (KVM) switch and one keyboard, monitor, and mouse for the entire cluster; alternately, one keyboard, monitor, and mouse per node
- Two Server Area Network (SAN) switches
- Two or more UPSes and power strips
- Cables for all of the above

This equipment can either sit on the floor or be assembled into one or more racks, as illustrated in AV 2-2 and AV 2-3 on page 2-6.

There are advantages and disadvantages to each solution, as listed in the table below:

Table 1: Rack and Tower Advantages and Disadvantages

Issue	Tower	Rack
Expense	Less expensive	More expensive
Cables	Harder to manage	Easier to manage
Weight Distribution	Weight is more distributed	Weight is more concentrated
Floor Space	Need more	Need less
Cooling	Less of a concern	More of a concern
Tilting injuries	Not an issue	An important issue

2.2 Cable Management

Regardless of the format you choose, the many components require a phenomenal number of cables. Cable management is a critical to the smooth assembly and operation of your cluster. When a cable or component needs to be replaced, you must be able to quickly and easily locate and replace the affected cables. Two factors are critical:

- Adequate space—you must allow sufficient space to maneuver around cables
- Labels—you need to be able to identify both ends of the cable, and the node to which it belongs

Label each cable for easy identification of its type, source, and destination. Color-coded labels help you identify cables by node. These conventions are in common use for NSC cluster hardware, but any code you use consistently will work as well.

Table 2: Cable Labeling Standard

Node Number	Cable Tie Color
1	Pink
2	Orange
3	Yellow
4	Green
5	Purple

Add a white tie to all the cables that connect on the X fabric. For example, put a white tie on all the cables that attach to node 1. If you have more than two nodes, use other colored ties to indicate which cables connect to which nodes.

2.3 Rack Operations

Racks allow you to concentrate equipment in a single vertical space. It is critical that you understand how to work with rack equipment before you begin loading and cabling the rack.

Some vendors supply planning guides or software to assist you in designing racks to accommodate your specific needs. Check the *NSC Curriculum Hardware Supplement* for more information about the use of racks.

2.3.1 Rack planning

Before loading any equipment into a rack, plan how the rack will be loaded:

- Consider equipment and rack sizes. Racks are supplied in varying depths and heights to accommodate different types and amounts of equipment. The height of racks and equipment is measured in an industry-standard units known as a *U*. Each *U* is 1.75 inches, or approximately 4.49 centimeters.

For example, if you had three nodes of 5 *Us* each, a monitor of 5*Us*, and a keyboard/mouse tray and KVM switch that each take 1 *U*, then you need at least a 22*U* rack to hold this equipment.

- Make sure the rack is deep enough to hold your largest component. Be sure to consider the depth of the rails as well.
- Consider whether you need open or closed racks.
 - Open racks are less expensive
 - Closed racks provide better ventilation.
 - Closed racks require slightly more assembly, because they require that you install doors and sides.
- Consider whether to use convective- or forced-air-cooled racks.
 - Convective racks are less expensive.
 - Forced-air-cooled racks provide better cooling.

- If you have more equipment than will fit into a single rack, use multiple racks. In this case:
 - The racks should be of the same size whenever possible.
 - Locate components so that power cables do not cross racks.
 - Each rack needs its own UPS.
 - Racks should bolt together to prevent tipping and cable abrasion.
- Consider future expansion.
- Make sure your room is tall enough for the rack(s) you choose.
- Make sure the floor is strong enough to hold the weight you plan to put into the rack.
- If the rack is pre-loaded, make sure it fits through all doorways without tipping. If a rack is too tall to fit through a door, you must unload it before tipping it.
- Consider mechanisms provided to prevent tipping. These include:
 - Pull-out extension, or permanent, legs in front
 - Side legs or feet

2.3.2 Rack assembly and loading


- Assemble the rack first. Always use the extension legs or stabilizer feet to prevent tipping, if they are provided with your rack.
- When loading the rack:
 - Assemble components as completely as possible, including rails, before you place the components into the rack
 - Keep the center of gravity as low as possible to avoid tipping the rack over:
 - Add components from the bottom up

- Put the heavier components closer to the bottom.


WARNING! A populated rack weighs enough to cause severe injury and possibly even death if it falls onto someone!

- Always place monitors at the top of a rack because they generate the most heat.
- If you must choose between weight and heat considerations, give weight highest priority, as you can also increase cooling capacity.
- Use at least two people to place heavier components in the rack. Because components can weigh hundreds of pounds, you need one person on each side of a component to lift the heavy components to the rack. If possible, have a third person guide the component into the rack and onto its rails. Always lift with your legs, never with your back.
- Do not extend more than one component at a time when removing equipment from rack.

AV 2-4



AV 2-5



2.3.3 Rack cabling

After all of the components are placed in the rack and the cables have all been labeled, you are ready to cable the components together.

The main issues when cabling are:

- Never run power cables between two racks. This creates a potential hazard and violates codes in many locales.
- Attach cables to the shortest components (depth-wise) first.
- Use cable-strain-relief ties where necessary. This also helps keep cables neatly bundled together.
- Use power strips designed for your rack between UPSes and other components. Place one power strip on each side of the cabinet, mid-height. Run power cables from the power strip to each power supply on the same side. This is shown in AV 2-4 on page 2-12.
- As shown in AV 2-5 on page 2-12, if using SAN switches or hubs, locate them centrally to the nodes.
- Be sure circuits are not overloaded if the rack is equipped with power strips.

WARNING! Do not overload the AC supply circuit that provides power to the rack to avoid the risk of fire, equipment damage, or personal injury.

2.4 Site Planning

Site planning is critical when deciding on a suitable location for the cluster equipment. Necessary attributes for a cluster site include (some of which have already been discussed):

Floor loading	Make sure the floor can handle the aggregate concentrated weight of the cluster equipment. This is especially important for clusters in racks, sites using raised floors, and sites using non-commercial class buildings.
Entry path	Make sure the cluster components can be moved from the delivery vehicle to the cluster site. • Can racks make it under doorways without being tipped? Again, if this is not the case, you must empty the rack before it is tipped. In addition, make sure doors are wide enough to fit the widest piece of equipment. • Are the hallways wide and tall enough? • Is the floor is smooth and level with no gaps?
Power	Make sure the circuits bringing power to the cluster have enough capacity.
Cabling	Make sure there is sufficient space around and between the hardware to allow for cable management. As discussed, the cabling requirements for clusters are substantial. It is critical that you leave an adequate amount of space to manage the cables in an organized fashion that allows you to quickly and easily swap cables and components.
Cooling	Make sure the cooling system for the location is adequate to handle the heat generated by the cluster hardware.

2.5 Lesson Summary

Some cluster configurations require a substantial amount of hardware. You might need to look for alternatives to traditional tower cabinet computers and storage devices. This lesson introduced rack-mount configurations and provided suggestions to help you plan for locating and assembling rack hardware.

Review questions

1. What are some advantages of using rack forms for cluster hardware?
2. What are the guidelines for labeling cables?
3. Describe the guidelines for placing components within and between racks.
4. What are some considerations for selecting a location for your cluster?

2.6 Answers to Exercise and Review Questions

This section contains the answers to review questions presented in this lesson.

Review Question Answers

This section contains the answers to the review questions on page 2-13.

1. Some advantages of using rack forms for cluster hardware are that they require less floor space and better facilitate cable management.
2. Label cables for easy identification of their type, source, and destination. Use color codes to indicate node membership. Add a white tie to uniquely identify the X fabric SAN cables.

3. Consider these guidelines for placing components within racks:

Assemble all components first, including rails

Place heaviest components at the bottom for stability

Place monitors at the top, for heat dissipation

Use two people to install heavy components

4. When planning a location for your cluster, consider power and cooling requirements, weight, cable space, and sizes of components.

3

Assembling and Initializing Cluster Hardware



Lesson 3

Assembling and Initializing Cluster Hardware

Introduction

You must follow a prescribed sequence of events in order to efficiently assemble a cluster and facilitate a successful software installation. Deviations from this established sequence will likely prove time-consuming and increase the likelihood that the software installation will not succeed.


This lesson details the rules for assembling a cluster.

Objectives

By the end of this lesson, you will be able to:

- Describe the general cluster hardware assembly process
- Assemble cluster hardware
- Assemble racks and rack-mounted equipment
- Initialize hardware to prepare for software installation

AV 3-1



Topics

In order to meet these objectives, this lesson discusses:

- High-level Overview — an overview of the assembly process
- Hardware Assembly — physically assembling these components of a cluster:
 - Racks
 - UPSes and power strips
 - Nodes and controllers — the various components and controller cards that need to be present in a node, whether using CNM or protected storage
 - Switches and hubs — placing and cabling switches and hubs for LAN, SAN, and keyboard/video/mouse (KVM) switches
 - Shared storage — placing and cabling the various shared storage methods
- Hardware Initialization — running low-level hardware initialization on the components of a cluster. The subsections for the components are:
 - Nodes — running CMOS setup on various kinds of nodes, including how to insure that attributes of certain controllers are set to specific values.
 - SAN — running diagnostics on the SAN fabrics
 - Shared storage — creating hardware RAID arrays on certain shared storage solutions.

The last part of this lesson is a lab where you assemble the hardware for either a 2-node cluster using cross-node mirroring, or a cluster with more than 2 nodes using shared storage.



3.1 Assembly Overview

This procedure will help you assemble your hardware in an orderly fashion. It has been tested and proven effective and complete.

If you have expertise assembling hardware, and racks in particular, you might choose to perform some tasks in a different order. While it is not absolutely required that you strictly follow these steps in the order given, you are strongly recommended to use this procedure because it is known to work. If you are not using racks, disregard rack-specific tasks.

1. Assemble racks, if used. See Section 3.2.1, “Assembling racks,” on page 3-11.
2. Prepare UPSes. Place them in the bottom of the racks, if used. See Section 3.2.2, “Assembling UPSes and power strips,” on page 3-13.
3. Assemble shared storage units, if used. The type of shared storage used dictates any required assembly. In general:
 - Place the rails on the storage units.
 - Load the storage unit into the rack.

Remember that racks must be loaded with the heaviest components toward the bottom. If the weight of the storage units, including disks, is more than some of your nodes, place them in the rack before those nodes. Otherwise, place them in the rack after those nodes.

- Place disks in the storage units after they are placed in the rack.
- See Section 3.2.3, “Assembling shared storage,” on page 3-15 for additional information.

4. Assemble each node:

- Insert any cards that are not pre-installed. For example, many systems do not come with NICs or HBAs.
- Insert SAN cards. Some SAN technologies requires you to insert SAN cards into particular slots.
- Place the rails on the nodes and place the nodes in the rack, heaviest nodes toward the bottom.

For more information, see Section 3.2.4, “Assembling nodes,” on page 3-17.

5. Install rack-mounted keyboard trays and place the tray in the rack at a comfortable height.
6. Assemble SAN switches or hubs. If using racks, attach rails and place them into the rack.
7. Assemble LAN switches or hubs. If using racks, attach rails and place them into the rack.
8. Place the KVM switch in the center of the nodes. If using a rack, attach rails to the KVM switch and install it in the rack.
9. Place the monitor in an appropriate location. If using a rack, install the monitor tray and place it in the rack.
10. Cable the components together, starting with the shortest components. For example, LAN, SAN, and KVM switches are short. Cable these components first.

Label each cable as described in Section 2.2, “Cable Management,” on page 2-8.

Note If you must install a longer component on top of a shorter component in a rack, you might want to attach cables to these shorter components before placing the longer component on top of it. This helps prevent abrasions on your fingers when cabling equipment in the rack.

11. Power on components in this order:

- a.) Monitors
- b.) Switches or hubs
- c.) External storage cabinets
- d.) Nodes

Make sure each component has completed its firmware initialization before powering on the next component. For example, verify that all storage cabinets have completed their firmware initialization before powering on any nodes.

12. Run any necessary firmware initialization on each node or storage cabinet. The order in which you must perform these steps is vendor-dependent. See Section 1.8, “Initializing Nodes,” in your vendor’s chapter in the ***NSC Curriculum Hardware Supplement***.
13. Run any necessary SAN verification software to verify that the SAN cards, cables, and switches or hubs are working properly.
14. Run any necessary LAN verification software to verify that the SAN cards, cables, and switches or hubs are working properly.

Once you have installed and initialized the hardware, you are ready to install the operating system as described in Lesson 4: “Installing SCO UnixWare 7.”

The next sections of this lesson provide greater detail and include references to vendor-specific information and procedures found in the ***NSC Curriculum Hardware Supplement***. The lab in Section 3.5, “Lab: Assembling Cluster Hardware,” on page 3-28 provides a complete procedure to use when assembling the classroom cluster hardware.

3.2 Hardware Assembly

Regardless of the vendor you chose for your cluster hardware, all have similar assembly procedures. The next sections describe common assembly tasks for cluster hardware, organized to reflect the order in which they are assembled and installed.

Tasks and information specific to each vendor's hardware components are described in the applicable section of the *NSC Curriculum Hardware Supplement*.

Use this checklist to insure that all components have been assembled and cabled before you continue to Section 3.3, "Hardware Initialization," on page 3-23.

Table 3: Assembly Checklist

Component Assembled	Component Cabled	Component Name
		Racks
		UPSes
		Power strips
		Nodes
		Shared storage
		SAN
		LAN
		KVM switch
		Monitor


3.2.1 Assembling racks

If you are unfamiliar with racks, see Section 2.3, “Rack Operations,” on page 2-9 before continuing. To assemble a rack:


1. Perform a site survey to find an appropriate location.
2. Select appropriate rack model(s).
3. Assemble the rack(s). When using closed racks, leave the doors and sides off until the components have been installed.
4. Load the equipment.
5. Ground the rack.

For details on assembling cluster-certified racks, see Section 1.4, “Assembling Racks,” in your vendor’s chapter in the *NSC Curriculum Hardware Supplement*.

AV 3-2



AV 3-3




3.2.2 Assembling UPSes and power strips

Although UPSes vary widely, they have certain things in common. For each UPS device:

1. Follow the manufacturer's specifications for connecting the battery.
2. Attach rails and place them in the rack, if used.
3. Run power cords from each UPS to the power source.
4. Plug the power strips into each UPS.

When you are finished configuring your UPSes and power strips, your configuration should look similar to those shown in AV 3-2 and AV 3-3 on page 3-12.

AV 3-4



3.2.3 Assembling shared storage

There are several shared storage solutions you can use with your cluster. Further, you can employ multiple solutions within a single cluster. For example, as shown in AV 3-4 on page 3-14, you may have potential root nodes connected to a fibre solution and two other nodes connected to a JBOD solution.

The supported shared storage options and their required hardware are discussed in the applicable sections of the Section 1.6, “Assembling Shared Storage,” in your vendor’s chapter in the *NSC Curriculum Hardware Supplement*.

To assemble shared storage:

1. Assemble the box(es).
2. Insert boxes into a rack, if used.
3. Populate the boxes with disks.
4. Cable the storage boxes to the nodes after nodes have been assembled.



3.2.4 Assembling nodes

When configuring a node:

1. Place HBAs in the node and connect them to appropriate disks.
 - You must have one HBA connected to one local disk that is used to boot the UnixWare 7 operating system on this node. This includes *boot*, *swap*, and *dump* slices.
 - If this node is a potential root node and you are using CNM, this will also be the disk used for the *root* filesystem.
 - If this node is protecting shared external storage, add one HBA for each external shared storage cabinet. There is an exception to this rule when using fibre storage with multiple fibre storage boxes attached to the node through fibre hubs.
2. Place SAN cards in the node. In many types of nodes, the bus and slot you choose for your SAN cards is critical. These are some common considerations for installing SAN cards:
 - Use the highest priority PCI slot
 - Use the fastest PCI bus
 - Do not put slower-speed PCI cards in the same bus with the SAN card; all cards then operate at the lowest-common speed

For example, Compaq ServerNet SAN cards have these requirements:

- Place the SAN card in the highest-priority PCI slot.
- Set the IRQ 15

If using other SAN technologies, install at least two SAN NICs per node.

See Section 1.5, “Assembling Nodes,” in your vendor’s chapter in the ***NSC Curriculum Hardware Supplement*** for details specific to your vendor.

3. Place LAN cards in the node. As with host machines, if the node is on multiple subnets, it must have one LAN NIC per subnet.
4. Attach rails to the nodes and place them in the rack, heaviest nodes toward the bottom.

AV 3-5

Compaq 8-port K/M/M Switch



Assembling05

© The Santa Cruz Operation, Inc.

3.2.5 Assembling SAN, LAN, and KVM switches or hubs

To assemble switches:

1. Assemble the SAN switches.

Recall that most SAN technologies require you to use a switch to connect more than two nodes in a cluster. Some switches require assembly and configuration. For information on assembling SAN switches, see Section 1.7, “Assembling SAN Switches,” in your vendor’s chapter in the *NSC Curriculum Hardware Supplement*.

2. Assemble LAN switches.

See your manufacturer’s documentation for information about installing and configuring LAN switches and hubs.

3. Cable the keyboard/video/mouse (KVM) switch, if used.

KVM switches are a single box, requiring no assembly. AV 3-5 on page 3-18 illustrates how to run cables between the nodes, the KVM, and the actual keyboard, mouse, and monitor in use.

Note Current KVM technology does not allow for redundant KVM switches. As such, the KVM can be viewed as a single point of failure for the cluster.

However, failure of the KVM can only cause nodes to fail if sufficient voltage is generated by the KVM failure. The damage to the node’s video, keyboard, or mouse circuitry caused by that voltage would have to be sufficient to subsequently damage the node itself.

This probability of this happening is so low that the KVM switch is not considered a potential single point of failure. If a KVM unit fails, it can be replaced without powering down the entire cluster.

4. Insert switches into the rack if used.

AV 3-6



3.2.6 Assembling monitors

To assemble a monitor:

1. Attach speakers to the monitor, if included.
2. Cable the speakers to the sound card in one of the nodes.
3. Sound card drivers are not yet clusterized and therefore represent a single point of failure, as do the speakers.
4. Attach the monitor base.
5. Place the monitor on a rack-mounted shelf, if used. As shown in AV 3-6 on page 3-20, monitors are usually placed at the top of a rack to allow heat to dissipate.
6. Attach power to the monitor.



3.3 Hardware Initialization

After you have assembled and cabled all of the hardware components, you must initialize certain components before they can be used in a cluster. This section discusses which components may need initialization, and how to perform that initialization. Where appropriate, you are referred to the *NSC Curriculum Hardware Supplement* for instructions specific to a particular make or model of a component.


3.3.1 Initializing nodes

Whether you need to initialize a node depends on these factors:

Node manufacturer	For example, all Compaq nodes must be initialized with SmartStart to insure proper ordering of the storage controllers.
SAN type	Some nodes require that you use a specific interrupt number and type for the SAN card. For example, Compaq ServerNet I cards must be set to interrupt at IRQ15 with a trigger-level-type interrupt signal. On some nodes, you control this through PCI Plug and Play BIOS configuration common to many computers. On others, you must use proprietary setup software.

See Section 1.8, “Initializing Nodes,” in your vendor’s chapter in the *NSC Curriculum Hardware Supplement* for information on setting up a particular make and model of node.

AV 3-7



3.3.2 Initializing shared storage

Most protected storage schemes use RAID mirroring, performed either by hardware or software. AV 3-7 on page 3-24 illustrates two examples:

- Cluster 1 uses a sophisticated fibre storage solution with built-in RAID. The box contains dual SCSI busses, controllers, and power supplies, so that a single storage box can be used without a single point of failure. In this scenario, the disks on one physical bus are mirrored with disks on the other bus. The box then presents each of the mirrors to the system as single disks.
- Cluster 2 uses a simple JBOD storage solution with no hardware RAID assistance. Since there is no redundancy in either SCSI busses or controllers, two storage boxes are required. In this case, the NSC software must establish a mirror between the disks on box 1 and those on box 2.

Hardware initialization of shared storage solutions is only necessary where hardware RAID is being used, as in the Cluster 1 example.

See Section 1.9, “Initializing Shared Storage,” in your vendor’s chapter in the ***NSC Curriculum Hardware Supplement*** for details on setting up specific shared storage solutions that utilize hardware RAID.

3.3.3 Installing the split-brain avoidance cable

Recall that a split-condition can exist when both nodes try to act as active root nodes. Certain SAN implementations that do not use SAN switches can experience this condition.

You can avoid this by configuring an alternate communication path using a null-modem cable between serial ports on two nodes. This prevents your *root* filesystem from being destroyed when the SAN communication is lost.

You are recommended to configure and test split-brain avoidance whenever you are installing a CNM cluster.

Attach a serial null-modem cable to a serial port on each node.

3.3.4 Testing the SAN

Before attempting to install the operating system software on a cluster, you must verify that the SAN components are functioning properly.

If using a vendor-specific SAN technology, such as Compaq ServerNet, see Section 1.11, “Testing the SAN Interconnect,” in your vendor’s chapter in the ***NSC Curriculum Hardware Supplement*** for information on testing you SAN components.

If you are using LAN technology to implement the SAN, test these components just as you would other LAN components. Verify:

- Card-to-card connectivity
- Switch or hub connectivity

3.3.5 LAN testing

Test your LAN NICs and hubs using software utilities provided by your LAN hardware vendor. Your instructor will provide any necessary test or diagnostic tools for your classroom LAN hardware.

3.4 Lesson Summary

This lesson, in addition to the *NSC Curriculum Hardware Supplement* presented information on how to assemble the hardware needed to implement an SCO NSC cluster, including:

- A step-by-step overview of the process
- Assembly and cabling of cluster components
- Initialization of cluster components

Review questions

1. If using a rack, which component would you assemble and install into the rack first, UPSes or nodes? Why?
2. Which component is unnecessary in a two-node cluster, but must be included on a similar cluster with more than two nodes?
3. When installing a cluster with more than two nodes, which protected storage methods can be used?
4. How does the installation of a cluster with CNM differ from one with external storage?
5. If you are using shared storage for your *root* filesystem, can all of your disks be in the external cabinets?

3.5 Lab: Assembling Cluster Hardware

This lab guides you through assembling a cluster with two or more nodes using one of the supported protected storage techniques. The instructor will provide you with a complete list of equipment being used in your classroom and the type of cluster you will assemble.

Perform one of the labs below, based on the protected storage solution you are using. During the lab, you must assemble any components not already assembled (for example, the rack). You must also initialize and test all applicable hardware components.

The labs are finished when you can show your implementation, including passing applicable hardware tests, to the instructor. Check your progress with the instructor at key phases of the lab, such as completion of SAN hardware initialization.

3.5.1 Cross-node mirroring procedure

In this lab, you will assemble a two-node cluster using cross-node mirroring between nodes 1 and 2 to protect the *root* filesystem.

This procedure is based on the instructions presented in Section 3.1, “Assembly Overview,” on page 3-7. Where appropriate, you will be referred to the ***NSC Curriculum Hardware Supplement*** for detailed information about the hardware in your classroom.

If your classroom does not use all of the hardware devices mentioned in these steps, such as racks, hubs, or switches, disregard those tasks.

- Step 1. Assemble the rack according to the manufacturer’s instructions. See Section 1.4, “Assembling Racks,” in your vendor’s chapter in the ***NSC Curriculum Hardware Supplement***.
- Step 2. Create a loading plan for the rack. See Section 1.4, “Assembling Racks,” in your vendor’s chapter in the ***NSC Curriculum Hardware Supplement*** for details about your manufacturer’s rack loading requirements.
- Step 3. Assemble SAN, LAN, and KVM cables according to the color-coding guide in Section 2.2, “Cable Management,” on page 2-8.
- Step 4. Assemble UPSes:
 - a. Connect their batteries, if necessary.
 - b. Place them in the rack, one pair in each rack.
 - c. If not using a rack, place the UPSes on the ground in a position central to where the rest of the components will reside.
 - d. Plug the UPSes into the wall outlets.
 - e. Mount the power strips on the side of each side of the rack(s).
 - f. Plug the power strips into the UPSes.

- Step 5. Assemble each node. For this CNM-based cluster, each node must have:
- One disk attached to an HBA. This disk will be used for *boot*, *swap*, *dump*, and *root* filesystems.
 - One or two SAN cards, depending on the SAN technology used. Remember, if using ServerNet I, this card must be placed in the highest priority PCI slot.
 - One LAN card

See Section 1.5, “Assembling Nodes,” in your vendor’s chapter in the *NSC Curriculum Hardware Supplement* for more detail about the nodes used in your classroom.

If using racks, attach rails to the nodes and place them in the rack according to the loading plan created in Step 1.

Note If installing SAN, LAN, or KVM switches between nodes, assemble but do not install nodes above the switches until the switches themselves have been installed. This will prevent injury to your fingers.

If using towers, place the nodes symmetrically on either side of the UPSes.

- Step 6. Assemble SAN switches if your cluster has more than two nodes. See Section 3.2.5, “Assembling SAN, LAN, and KVM switches or hubs,” on page 3-19.

If using racks, attach rails to the switches and place them in the rack.

Note After installing the SAN switch, but before installing other devices above it, connect the SAN cables to the back of the SAN switch. Then install other devices above the SAN switch. Attaching cables to the switch at this time makes it easier to see and access the connection points.

If using a tower system, place the switches as centrally as possible to the nodes. A common location for the switches in tower clusters is on top of the nodes themselves.

Step 7. Assemble LAN switches or hubs, if used.

- Place them in the rack, if used. Attach cables before adding any devices above them in the rack.
- If using a tower system, place the LAN switches or hubs on top of the SAN switches or hubs. Attach cables.

Step 8. Assemble KVM switches, if used.

- If using a rack, attach rails to the KVM switch and install it in the rack. Attach cables to the KVM switch before adding any devices above it in the rack.
- If using a tower system, place the KVM switch on top of the network switches or hubs.

Step 9. Connect the keyboard.

Install the keyboard tray, if using a rack. Position it at a comfortable height to allow the average person to type while standing in front of the cluster.

Note Install remaining nodes within the rack above the keyboard tray and below the monitor mentioned next.
--

If you are not using a rack, place the keyboard on top of one of the nodes.

Step 10. Connect the monitor.

If using a rack and a monitor tray within the rack, place the monitor tray at the top of the rack. For racks less than full height, you can place the monitor on top of the rack itself.

Step 11. Connect the LAN cables from the LAN hub or switch to the LAN cards in each node. Use the color-coded labels as a guide.

Step 12. Connect the SAN cables. See Section 1.7, “Assembling SAN Switches,” in your vendor’s chapter in the *NSC Curriculum Hardware Supplement* for instructions for your classroom equipment.

Use the color coding to decide which cable to plug into which node.

- Step 13. Connect the KVM cables from the KVM switch to the keyboard, video, and mouse connectors on each node. Use the color coding to decide which cable to plug into which node.
- Step 14. Install the split-brain avoidance cable, if using a CNM cluster.
- Step 15. Power on the monitor.
- Step 16. Power on any switches or hubs.
- Step 17. Power on node 1.
- Step 18. Initialize the node. See Section 1.8, “Initializing Nodes,” in your vendor’s chapter in the *NSC Curriculum Hardware Supplement* for instructions specific to your classroom hardware.
- Step 19. Repeat Steps 15 and 16 for each node in the cluster.
- Step 20. Perform LAN diagnostics using software passed out by your instructor.
- Step 21. Perform SAN diagnostics. See Section 1.11, “Testing the SAN Interconnect,” in your vendor’s chapter in the *NSC Curriculum Hardware Supplement* for instructions specific to the SAN hardware in your classroom. Your instructor will provide any necessary software.

Your cluster hardware is now ready to install UnixWare 7 operating system software.

3.5.2 Fibre Channel/JBOD/RSO procedure

In this lab, you will assemble a two (or more)-node cluster using fibre channel between nodes 1 and 2 to protect the *root* filesystem.

These instructions are based on those presented in Section 3.1, “Assembly Overview,” on page 3-7. Where appropriate, you will be referred to the ***NSC Curriculum Hardware Supplement*** for a detailed list of hardware make and model specific instructions.

- Step 1. Assemble the rack according to the manufacturer’s instructions. See Section 1.4, “Assembling Racks,” in your vendor’s chapter in the ***NSC Curriculum Hardware Supplement***.
- Step 2. Create a loading plan for the rack. See Section 3.2.1, “Assembling racks,” on page 3-11.
- Step 3. Label the SAN, LAN, and KVM cables according to the color-coding guide in Section 2.2, “Cable Management,” on page 2-8.
- Step 4. Assemble UPSes:
- Connect their batteries, if necessary.
 - Place them in the rack, one pair in each rack.
 - If not using a rack, place the UPSes on the ground in a position central to where the rest of the components will reside.
 - Plug the UPSes into the wall outlets
 - Mount the power strips on the side of each side of the rack(s)
 - Plug the power strips into the UPSes
- Step 5. Assemble the shared storage unit(s). See Section 1.6, “Assembling Shared Storage,” in your vendor’s chapter in the ***NSC Curriculum Hardware Supplement***.

If you are using a tower system, locate the shared storage box(es) wherever it is convenient, so long as the fibre cables can reach between the storage box and its nodes.

If you are using a rack, place the shared storage box(es) in the rack according to your plan. However, do not populate the box(es) until they have been placed in the rack.

Step 6. Assemble each node. Since you are using external storage units for protected storage, each node must have:

- One internal disk attached to an HBA. This disk will be used for boot, swap, dump. It will not contain the root filesystem.
- One or more HBAs to attach to the storage box(es).
- One or two SAN cards, depending on the SAN technology used. Remember, if using Compaq ServerNet I, this card must be placed in the PCI slot with the highest priority.
- One LAN card

See Section 1.5, “Assembling Nodes,” in your vendor’s chapter in the ***NSC Curriculum Hardware Supplement*** for details on assembling your specific nodes.

If using racks, attach rails to the nodes and place them in the rack according to the loading plan created in Step 1.

Note If installing SAN, LAN, or KVM switches between nodes, assemble but do not install nodes above the switches until the switches themselves have been installed. This will prevent injury to your fingers.

If using towers, place the nodes symmetrically on either side of the UPSes.

Step 7. Assemble SAN switches if your cluster has more than two nodes. See Section 1.7, “Assembling SAN Switches,” in your vendor’s chapter in the ***NSC Curriculum Hardware Supplement***.

If using racks, attach rails to the switches and place them in the rack.

Note After installing the SAN switch, but before installing other devices above it, connect the SAN cables to the back of the SAN switch. Then install other devices above the SAN switch. Attaching cables to the switch at this time makes it easier to see and access the connection points.

If using a tower system, place the switches as centrally as possible to the nodes. A common location for the switches in tower clusters is on top of the nodes themselves.

- Step 8. Assemble LAN switches or hubs, if used.
- Place them in the rack, if used. Attach cables before adding any devices above them in the rack.
 - If using a tower system, place the LAN switches or hubs on top of the SAN switches or hubs. Attach cables.

- Step 9. Assemble KVM switches, if used.
- If using a rack, attach rails to the KVM switch and install it in the rack. Attach cables to the KVM switch before adding any devices above it in the rack.
 - If using a tower system, place the KVM switch on top of the network switches or hubs.

- Step 10. Connect the keyboard.
- Install the keyboard tray, if using a rack. Position it at a comfortable height to allow the average person to type while standing in front of the cluster.

Note Install remaining nodes within the rack above the keyboard tray and below the monitor mentioned next.
--

If you are not using a rack, place the keyboard on top of one of the nodes.

- Step 11. Connect the monitor.
- If using a rack and a monitor tray within the rack, place the monitor tray at the top of the rack. For racks less than full height, you can place the monitor on top of the rack itself.

- Step 12. Connect the LAN cables from the LAN hub or switch to the LAN cards in each node. Use the color-coded labels as a guide.

- Step 13. Connect the SAN cables.
- Use the color coding to decide which cable to plug into which node.

- Step 14. Connect the KVM cables from the KVM switch to the keyboard, video, and mouse connectors on each node. Use the color coding to decide which cable to plug into which node.

- Step 15. Cable the external shared storage unit, if used. See Section 1.6, “Assembling Shared Storage,” in your vendor’s chapter in the ***NSC Curriculum Hardware Supplement***.
- Step 16. Power on the monitor.
- Step 17. Power on any switches or hubs.
- Step 18. Power on node 1.
- Step 19. Initialize the node. See Section 1.8, “Initializing Nodes,” in your vendor’s chapter in the ***NSC Curriculum Hardware Supplement*** for instructions specific to your classroom hardware.
- Step 20. Repeat Steps 17 and 18 for each node in the cluster.
- Step 21. Initialize the disks and controllers in the shared storage units. See Section 1.9, “Initializing Shared Storage,” in your vendor’s chapter in the ***NSC Curriculum Hardware Supplement*** for instructions specific to your classroom hardware.

Note Depending on the Fibre shared storage solution used, you might need to perform additional initialization tasks on node 2. If so, these steps are included in the applicable section of the <i>NSC Curriculum Hardware Supplement</i>.

- Step 22. Perform LAN diagnostics using software passed out by your instructor.
- Step 23. Perform SAN diagnostics. See Section 1.11, “Testing the SAN Interconnect,” in your vendor’s chapter in the ***NSC Curriculum Hardware Supplement***. Your instructor will provide any necessary software.

Your cluster hardware is now ready to install UnixWare 7 operating system software.

3.6 Answers to Review Questions

This section contains the answers to review questions presented in this lesson.

Review Question Answers

This section contains the answers to the review questions on page 3-27.

1. According to the procedure on Section 3.1, “Assembly Overview,” on page 3-7 and the explanation given in Section 3.2.2, “Assembling UPSes and power strips,” on page 3-13, you should assemble and install UPSes before nodes because UPSes weigh more per U than nodes, and such items must be placed lower in the rack to keep the rack’s center of gravity low.
2. According to Section 1.7, “Assembling SAN Switches,” in your vendor’s chapter in the *NSC Curriculum Hardware Supplement*, SAN switches are only required for clusters with more than two nodes. As such, in a cluster with only two nodes, said switch need not be included.
3. According to Section 1.2, “Supported Protected Storage Methods,” in your vendor’s chapter in the *NSC Curriculum Hardware Supplement*, the protected storage method you use has nothing to do with the number of nodes in a cluster. As such, all supported protected storage methods are available on the cluster in question.
4. According to Section 3.2.4, “Assembling nodes,” on page 3-17, CNM systems only need HBAs for their internal disks, which will serve as both boot and root disks. For non-CNM systems, you must have an additional HBA per shared storage box connected to the node.
5. According to Section 3.2.4, “Assembling nodes,” on page 3-17, you must have a boot disk in addition to the external root disk for systems not using CNM.



4

Installing SCO UnixWare 7



Lesson 4

Installing SCO UnixWare 7

Introduction

Once you have assembled and initialized your cluster hardware, install UnixWare 7 on the initial root node.

To set up the proper environment for the subsequent installation of the NSC software, you must carefully navigate through the UnixWare 7 installation procedure, taking into account the NSC-specific considerations described in this lesson.


The lesson focuses on the unique NSC properties that influence the UnixWare 7.1 installation process.

Objectives

By the end of this lesson, you will be able to:

- Describe the installation process
- Prepare for the selections you must make during UnixWare 7 installation to support NSC functionality
- Perform any necessary pre-install tasks required by your cluster hardware vendor
- Install UnixWare 7 on the initial root node
- Perform post-installation tasks to prepare for NSC software installation

AV 4-1



Topics

In order to meet these objectives, this lesson discusses:

- Vendor-supplied EFSes — checking for hardware-specific drivers needed for installation
- Licensing — selecting appropriate UnixWare license editions for NSC
- Cluster Virtual Host Name — selecting appropriate node and cluster host names
- Disk and Filesystems — special considerations for NSC supported functionality
- Package Selection — deciding which packages are required, and which should be removed

4.1 Understanding Software Installation

After your hardware is assembled and initialized, your cluster software is assembled and installed in this sequence. This lesson covers Steps 1, 2, and 3.

1. Assemble operating system software. Create or obtain vendor-supplied EFS or HBA diskettes, if needed.
2. Install UnixWare 7.
3. Install vendor-supplied EFSes.
4. Prepare the initial root node for NSC installation
5. Install NSC on the initial root node
6. Install the failover root node
7. Install depending nodes.

Successful installation of a cluster requires careful preparation. Because installation of all nodes depends on the success of the installation and preparation of the initial root node, you must understand the configuration choices you will make during the UnixWare 7 installation process that are needed to support the NSC software.

This lesson describes steps 1, 2, and 3 of the software installation procedure. Each section gives you guidance about the selections you need to make during the UnixWare 7 installation process that affect the NSC software preparation and installation steps.

The lab in Section 4.10, “Lab: Installing UnixWare 7.1 on the Initial Root Node,” on page 4-20, provides the detailed procedure. Lesson 5: “Installing NSC Software” and Lesson 6: “Installing the Failover Root Node” describe the procedures for installing the cluster software.

4.2 Special HBA driver considerations

Hardware vendors supply HBA drivers for devices that are not included in the UnixWare 7.1 Media Kit. These are supplied in a number of ways:

- HBA diskettes supplied by the hardware manufacturer
- CD-ROMs provided by hardware vendors with HBA diskette images
- EFS diskettes sets which include one or more HBA diskettes
- Downloadable HBA diskette images

Be sure to read Section 1.12, “Vendor-Supplied Software,” in your vendor’s chapter in the ***NSC Curriculum Hardware Supplement*** and use only those drivers specified by your hardware vendor. Be careful to install HBA disks in the order specified.

4.3 Licensing Information

There are two types of software licenses needed for a cluster:

- One UnixWare 7.1 license per node
- One NonStop Clusters license

Be sure you understand the license requirements and have the appropriate license information from the Certification of License Authority (COLA) card before beginning your UnixWare 7 installation.

You might also have other licenses for optional products or services.

4.3.1 UnixWare 7 licenses

Each node in the cluster must have a UnixWare 7 license. SCO provides licenses separately from the installation media. The UnixWare 7 Media Kit contains the software installation CD-ROM disks. A separate license package includes a COLA for each node with this information:

- License number
- License code
- License data

Each node must be installed with the same license UnixWare 7 *Edition*. Editions are collections of software targeted to a specific purpose. You must install each node with an Edition that includes a Veritas Online Data Manager (ODM) or Disk Mirroring license.

Alternately, you can add an optional ODM or Disk Mirroring license to an Edition that does not include mirroring or ODM.

ODM and Disk Mirroring licenses are included with these editions:

- Enterprise
- Data Center

An optional ODM or Disk Mirroring license can be purchased and used with these editions:

- Base
- Business
- Department

You cannot use the Free Unix, or the 60 Day Evaluation Editions because they do not support ODM.

If you install UnixWare 7 with a license Edition that does not include ODM or Disk Mirroring, you can add one of these licenses after UnixWare 7 installation using the **License Manager**. However, this *must* be added *before* NSC software is installed.

4.3.2 NSC licenses

In addition to the UnixWare 7 licenses required for each node, a separate NSC license is needed for the cluster software. At a minimum, a NonStop Clusters 2-Server Base License is required. If you add nodes, you must purchase NonStop Clusters Additional Server Licenses.

Note Do not enter your NonStop Clusters license during the UnixWare 7 operating system installation procedure.

You can install the NonStop Clusters software without a license for evaluation purposes. After 60 days, you must purchase and install NonStop Cluster licenses or remove the software from each node.

4.3.3 Other licenses

You might require additional licenses to support your specific cluster configuration. Some licenses apply to specific nodes, and some to the entire cluster.

For example, if you need to support 200 simultaneous users accessing the cluster, you must purchase additional Network User licenses. User licenses are applied to the cluster.

You might also require licenses for multiprocessor nodes, or to support nodes with large amounts of memory. CPU and RAM licenses apply to a single node.

All additional licenses should be installed using the **License Manager** after NSC is installed and all nodes are clusterized.

Consult the UnixWare 7 Installation Roadmap document supplied with the UnixWare 7 Media Kit to determine which licenses you need for your cluster.

4.4 Cluster Virtual Host Name

During UnixWare 7 installation, you are prompted to enter a node name. For a typical installation of a standalone server, this is the name by which the system is known locally and on the network. The commands **uname(1)** and **hostname(1)** return this name.

When installing the initial root cluster node, the node name you enter becomes the cluster virtual host name; that is, the collective name for all the nodes in the cluster. Therefore, you must supply the name that you want to call the cluster when you install UnixWare 7 on the initial root node.

Later, when you install the NSC software, you are prompted to enter the node name for the server itself.

4.5 Understanding Disk and Filesystem Configuration

NSC has unique requirements when configuring the filesystems. Be sure you understand these cluster-specific considerations before you start installing UnixWare 7 on the initial root node.

- Set */tmp*, */var*, and */var/tmp* to *off*. They are set to *memfs* by default, which only exist in the node's RAM and is therefore not available to other nodes. When set to *off*, these are created as directories within the *root* filesystem, and so they are *clusterized* and available to all nodes.
- Increase */stand* to a minimum of 40 MB. Additional copies of the kernel are stored here.
- Set up your *root* filesystem as recommended in Section 4.5.1, "The cluster root filesystem," on page 4-13.
- Do not create user filesystems such as */home*. Add these after installing NSC. Note that if you create user filesystems, they are not automatically *protected*; that is, mirrored.
- Add the values in the size column and verify that the space allocated is less than or equal to the size of the disk.

This table shows an example of the filesystem configuration table for an initial root node.

Table 4: Recommended Filesystem Setup

Name	Type	Size	Disk
/	vxfs	User configuration; minimum 2000, recommended 3000. Allow room for OSU	1
/stand	bfs	40	1
/dev/swap	slice	Increase calculated value	1
/dev/dump	slice	Calculated based on installed memory	1
/home	off	Do not create until after installing NSC	-
/home2	off	Do not create until after installing NSC	-
/var	off	-	-
/tmp	off	-	-
/var/tmp	off	-	-
/dev/volprivate	slice	Calculated; do not change	1
ALTS TABLE	slice	Calculated; do not change	1

4.5.1 The cluster *root* filesystem

The initial root node is the system on which you install UnixWare 7 and then perform the first installation of SCO NonStop Clusters software. Carefully consider how to configure your *root* filesystem before starting installation of the initial root node by familiarizing yourself with the unique aspects of the *root* filesystem in a cluster.

The *root* filesystem in a cluster is controlled by the initial root node, sometimes called the initial installation node. In the event of root node failure, the failover root node assumes control of the root filesystem.

- In a shared disk configuration, the *root* filesystem resides on an external disk shared by the root and failover root nodes.
- In a cross-node mirroring configuration, the *root* filesystem resides on the internal disk of the root node and is mirrored on the internal disk of the failover node.

The rest of this section provides guidelines for planning the size of the root filesystem to best balance your system requirements with the configuration of your cluster.

4.5.1.1 Planning *root* filesystem size

During installation, you are required to specify the size of the *root* filesystem. Calculate the size carefully. A very large *root* filesystem can significantly extend installation time. When determining the *root* filesystem size, different considerations apply to shared-disk clusters and others to those using cross-node mirroring.

Shared disk The NSC installation initially sets up the *root* filesystem on the internal drive of the initial installation node. Later during the installation, it moves the *root* filesystem to the external drive. Therefore, you cannot specify a *root* filesystem size that is larger than the internal disk of the initial installation node.

Balance operational needs and the time it takes for the *root* filesystem to be moved. Here are some examples:

- On systems where the initial installation node has a 9G internal drive, the *root* filesystem size should be a minimum of 2G.
- An average *root* filesystem size is around 3G.
- The *root* filesystem on node 1 must be smaller than the size of the drive in the external storage unit that will host the *root* filesystem.

If you plan to configure your system for future Online Software Upgrade (OSU) service, the drive that will eventually contain your *root* filesystem:

- Must be at least twice the size of the *root* filesystem
- Must also have room for any non-*root* filesystems you plan to store there

For example, if you have a 9G drive in your initial installation node and a 9G drive in your external array drive, and you plan to use OSU, the largest *root* filesystem you could specify would be limited to 4.5G in size; however, that would leave no room for any non-*root* filesystems on the external drive.

CNM The internal disk of the initial root node and that of the failover root node must each hold the *root* filesystem (and a spare copy of that filesystem if you plan to configure for OSU usage), as well as any non-*root* filesystems.

The *root* filesystem on node 1 must be smaller than the size of the internal drive in node 2.

4.5.2 Online System Upgrade (OSU)

During installation, you are prompted to set aside space for future Online System Upgrade (OSU) use. The OSU facility allows you to update UnixWare and NSC software while the cluster is fully-functional.

When you plan current filesystem space on the drive containing the *root* filesystem, include enough space for future OSU use. You must have room on the drive containing the *root* filesystem for a duplicate of that *root* filesystem.

Note Not all versions of NSC support OSU for all types of protected storage methods. Check the <i>SCO NonStop Clusters Release Notes</i>.
--

If you have additional non-*root* filesystems on the drive containing the *root* filesystem:

- Subtract the non-*root* filesystem sizes from the total space on the drive.
- The remaining space can be used for the *root* filesystem and the OSU set-aside space.
- Make sure your *root* filesystem is half the remaining space or less.

For example, if you have an 8GB drive, and non-*root* filesystems take up 1GB of the drive, you have 7GB of remaining space to use for the *root* filesystems. To enable the OSU feature, you should make the *root* filesystem no larger than 3.5GB.

4.5.3 Planning non-*root* filesystems

You can use any unused space on cross-node mirroring root volumes, or on shared-root external and internal drives for non-*root* files systems, such as */home* or */home2*. How you create non-*root* filesystems depends on the type of protected shared storage you use.

To simplify installation, do not configure any non-*root* filesystems until after UnixWare 7 and NSC have been installed and tested.

4.5.4 Vendor-specific partitions

Some vendors use a portion of the boot disk to store hardware- or vendor-specific utilities. Be sure to leave special partitions intact during the UnixWare 7 installation.

4.6 Deferring Network Configuration

Network configuration of the initial root node must be done in a particular sequence that is unique to NSC—after UnixWare 7 and supplements, but before NSC software.

Network configuration is deferred until after UnixWare 7.1 is running because there might be updated drivers that are installed later. For example, your hardware vendor might supply an EFS with NIC drivers that must be installed after UnixWare 7 is running.

The NIC must be configured before NSC because the IP address you assign becomes the CVIP during NSC installation. Likewise, the host name you configure for the first network interface becomes the cluster host name.

4.7 Package Selection

It is crucial that you select an NSC-compatible set of packages when installing UnixWare 7 on the initial root cluster node. Deselecting some packages can cause the NSC installation to fail. Conversely, installing some packages will cause problems on an NSC cluster node. It is strongly recommended that you adhere strictly to these guidelines:

- Select these packages from the UnixWare 7 Installation CD-ROM (1 of 3):
 - OS Multiprocessor Support—required *even* if you do not have multiprocessor server
 - Kernel Debugger
- Remove these packages from the UnixWare 7 Updates CD-ROM (2 of 3), if they are pre-selected:
 - Multipath IO (MPIO)
 - SCO NetWare
 - NIS
- Make sure the ODM package selection is FULL
- Defer installing the third UnixWare 7 CD-ROM until you have successfully completed the operating system installation.

Detailed instructions for making package selections are provided in Section 4.10, “Lab: Installing UnixWare 7.1 on the Initial Root Node,” on page 4-20.

4.8 Using EFSes and PTFs

Your hardware vendor provides Enhanced Functionality Supplements (EFSes) with the UnixWare 7.1 drivers you need to support your cluster hardware. These can be provided in several ways. For example:

- Floppy disk images created from a CD-ROM
- Floppy disk images downloaded from a suppliers web site

You might have a set of EFSes required to initially install UnixWare 7.1, and another to provide enhanced NSC-specific features. Refer to Section 1.12, “Vendor-Supplied Software,” in your vendor’s chapter in the *NSC Curriculum Hardware Supplement*. Be sure to note whether vendor supplements are required to be installed in a particular order.

During the installation procedure provided in the Lab, you are directed to sections in the *NSC Curriculum Hardware Supplement* that contain specific information for the hardware used in your classroom.

4.8.1 Program Temporary Fixes (PTFs)

SCO and your hardware vendor provide Program Temporary Fixes (PTFs) which are applicable to specific versions of UnixWare 7 and NonStop Clusters.

Use only those PTFs specified in the installation procedures in Lesson 5: “Installing NSC Software.” Non-clusterized PTFs are not supported on UnixWare 7 systems running NSC.

Check the *SCO NonStop Clusters Release Notes* for guidance about using PTFs.

4.9 Lesson Summary

This lesson provided guidelines for installing UnixWare 7 on your initial root cluster node. The selections that impact the subsequent NSC installation process were highlighted and where necessary, specific directions were provided to help you navigate through critical sections of the installation process.

Review questions

1. In which order are HBA disks usually inserted during UnixWare 7 installation?
2. How much space should you allocate for the root filesystem if you have a 4 GB UNIX partition on the internal primary disk drive and want to use the OSU option?
3. Can you use Department Edition UnixWare 7 licenses on your cluster nodes?
4. If you create a */home* filesystem during UnixWare 7 installation, will it automatically be mirrored to the failover node when NSC is installed?

4.10 Lab: Installing UnixWare 7.1 on the Initial Root Node

This lab guides you through the UnixWare 7.1 installation for the initial root node of the classroom cluster to which you are assigned. The instructor will provide you with any configuration information you need.

The steps in the procedure are general with references to hardware- or vendor-specific sections of the *NSC Curriculum Hardware Supplement*. Be sure to read through the procedure and the referenced hardware-related material before beginning the lab.

The lab is finished when you can show your functional UnixWare 7.1 root node is running properly and connected to the classroom network. Check your progress with the instructor at key phases of the lab.

This lab assumes you have performed all hardware assembly and initialization labs in this course.

Procedure

Use this procedure to install UnixWare 7. The hardware- or vendor-specific information you need is included in the NSC Hardware Supplement. When necessary, these steps refer you to the application section of the Supplement.

Recall that UnixWare 7 installation is divided into several stages:

Setup	The installation prompts you with questions and records your answers. Your replies determine which software components are installed, how system hardware is configured, and the disk division layout.
Software loading	The installation program installs and configures your system using the parameters defined during setup. This stage is not interactive.
Initial hard disk boot	Configuring and testing your mouse.
Post-install setup	Adding EFSes and application software.

To install UnixWare 7 on the initial node:

- Step 1. Create any HBA disks required for your classroom hardware. See Section 1.12, “Vendor-Supplied Software,” in your vendor’s chapter in the ***NSC Curriculum Hardware Supplement***.
- Step 2. Power off all nodes and external storage units using the procedure recommended by your hardware vendor. See Section 1.10, “Powering Down Cluster Hardware,” in your vendor’s chapter in the ***NSC Curriculum Hardware Supplement***.
- Step 3. Insert the first UnixWare 7 installation diskette into the primary diskette drive of the initial root node.
- Step 4. Power on the node.

You see your system’s boot messages, and then SCO logo. It might take several minutes to get to this point.
- Step 5. Set the language, if prompted.
- Step 6. Insert the second UnixWare 7 installation diskette when prompted.

After software is extracted from this diskette, the **Welcome** window appears. Press **F10** to continue.

Step 7. Select the zone, locale, keyboard.

Step 8. Enter the UnixWare 7 license number, code, and data for node 1. See Section 4.3, “Licensing Information,” on page 4-8, for help understanding the license requirements for a cluster node.

Note that the NSC license is *not* entered here.

Step 9. Install Host Bus Adapter disk(s) when prompted.

- If you have a vendor-supplied HBA disk(s), read Section 1.12, “Vendor-Supplied Software,” in your vendor’s chapter in the **NSC Curriculum Hardware Supplement** to make sure you insert the HBA disk(s) in the required order.
- You must insert the SCO HBA diskette in addition to any vendor-supplied HBA diskette(s). You typically install vendor-supplied HBAs first.
- If you do not have a vendor-specific HBA disk, insert the SCO UnixWare 7 HBA disk.

Select “Proceed with Installation” and press **F10** after you have entered all necessary HBA diskettes.

Step 10. Select “Do not enter the DCU (autoconfigure drivers)” when prompted.

Step 11. Enter the name of the cluster when prompted for the system node name.

Note Be sure you enter a name by which the entire cluster will be known. You will enter an actual system node name for this server later when installing NSC.

See Section 4.4, “Cluster Virtual Host Name,” on page 4-11 for more information.

Step 12. Select “Install from CD-ROM” and press **F10**.

Step 13. Insert the first UnixWare 7 CD-ROM (1 of 3) when prompted and press **F10**.

Step 14. Select “Customize partitions” in the “Action” field for disk number 1. Press **F10** to continue.

The partition table for disk 1 is displayed. You might already have a partition on your primary disk. .

Note Some hardware vendors create a disk partition to store configuration utilities. Do not remove this partition

- Step 15. Add a UNIX partition. Scroll to an unused partition number and setting the disk options. Use **→** to move between fields and **F2** to select these values:
- Choose “UNIX System” in the “Type” field.
 - Choose “Active” in the “Status” field.
 - Use the remainder of the primary disk for the new UNIX partition.
 - Press **F10** to continue.
- Step 16. Configure additional internal disks to use the entire disk for UNIX, if present in your node.
- Step 17. Select “Customize filesystems and slices” and press **F10**.
- Make these changes:
- Set */tmp*, */var*, and */var/tmp* to off.
 - Increase */stand* to at least 40 MB.
 - Increase */swap*. It is automatically decreased when you turn off */tmp*, */var*, and */var/tmp*.
 - Set your *root* filesystem size according to the guidelines in Section 4.5.1, “The cluster root filesystem,” on page 4-13.
 - Do not create user filesystems such as */home* but leave space for adding them later, after installing NSC.
 - Add the values in the size column and verify that the space allocated is less than or equal to the size of the disk.
 - When finished, press **F10**.
- Step 18. Select “Customize disk options” and press **F10**.

- Step 19. Mark all choices “no” *except* “Install a new boot sector.” If not already set, change “Install a new boot sector” to “yes” and press **F10**.
- Step 20. Select “Customize Installation of Packages” when prompted to choose a system profile and press **F10**

WARNING! Do not deviate from the package selections documented here. This can cause your NSC installation to fail later.

- Step 21. Select “Core System Services” and press **F10**.
- Step 22. Select “Extended OS Utilities” and press **F2**.
- Step 23. Select “Select individual packages” and press **ENTER**. Select “Kernel Debugger.”
- Step 24. Page down and select “OS Multiprocessor Support (OSMP).”

You MUST install OSMP even if you have a single-processor system.

Note Do NOT select MPIO.

- Step 25. Press **ENTER** when finished with Extended OS Utilities selection.
- Step 26. Select “Networking Services” and press **F10**.
- Step 27. De-select “NIS” and press **F2**.
- Step 28. Page down and select “SCO Netware.” Press **F2** and select the option “Select no packages (OFF).” Press **ENTER**.
- Step 29. Verify that the “Online Virtual Disk Volume Administration” selection is set to “FULL” in the **Select/De-select Services** window.
- If it is not, you do not have an appropriate UnixWare 7 license. You can continue installation, but must add an ODM or Mirroring license after UnixWare 7 is install, but before NSC is installed.
- Step 30. Page down and select “SCO Netware.” Press **F2** and select the option “Select no packages (OFF).” Press **ENTER**.
- Step 31. Press **F10** when finished making these selections.

- Step 32. Press  in the **Current Selections** window to continue.
- Step 33. Select “Intel MP Specification” as the Platform Support Module *even if you do not have a multiprocessor system*. Press .
- Step 34. Select “Defer Network Configuration” and press .
- Step 35. Set the date and time and press .
- Step 36. Select a security level and press .
- Step 37. Enter the system owner information and press .
- Step 38. Set the **root** password and press .
- Step 39. Accept the license agreement and press .
- Step 40. Press  to start installing the software.
- Software loading is the longest stage of the installation and can proceed unattended. You can track the status of the installation by viewing the progress indicator.
- Step 41. Reinsert your UnixWare 7 HBA diskette when prompted.
- Step 42. Reinsert any vendor-supplied HBA diskette(s).
- Step 43. Remove your floppy disk and CD-ROM when the installation completes.
- Step 44. Press  to shut down and restart your system.

4.10.1 Continuing with the installation

After the restart of your system and several non-interactive configurations take place, you are prompted to configure your mouse. Follow these steps to configure your mouse.

- Step 1. Select the type of mouse. Press **F10**.
- Step 2. Confirm or change the mouse parameters. Press **F10**.
- Step 3. Press **ENTER** to begin the fifteen (15) second mouse test.
- Step 4. Click on your right mouse button. Press **ENTER** to continue.
- Step 5. Insert the UnixWare Upgrades CD-ROM (2 of 3) and press **F10**.
- Step 6. Scroll to “Enhanced Event Logging System” and press **SPACE** to deselect it. Press **ENTER** to continue.

Software is now loaded from this CD-ROM. You can track the status of the installation by viewing the progress indicator.

Press **F10** when complete, and remove the CD-ROM.

- Step 7. Press **F8** to defer installing the Optional Services CD-ROM (3 of 3). Wait while the kernel is rebuilt and the system reboots.

Your initial root node is now running UnixWare 7.

4.10.2 Post-installation tasks

After UnixWare 7 has been installed on your initial root node, and before you install NSC software:

- Step 1. Add an ODM or Mirror license, if the license Edition used during UnixWare 7 installation did not include one of these licenses.
- Step 2. Install any vendor-supplied EFSes or applications. See Section 1.12, “Vendor-Supplied Software,” in your vendor’s chapter in the ***NSC Curriculum Hardware Supplement***.
- Step 3. Install all application software. It is unlikely you will do this in the classroom.
- Step 4. Reboot your system.

4.11 Answers to Exercise and Review Questions

This section contains the answers to review questions presented in this lesson.

Review Question Answers

This section contains the answers to the review questions on page 4-19.

1. EFSes supplied by hardware manufacturers are usually installed first. Then the HBA diskette supplied by SCO with the UnixWare 7 Media Kit is installed.
2. The *root* filesystem can be no larger than 2 GB since you must have space for two complete *root* filesystems in the UNIX partition.
3. You can use a Department Edition UnixWare 7 licenses for cluster nodes. However, you must add an ODM or Disk Mirroring license.
4. If you create a */home* filesystem it is not automatically configured as a mirrored filesystem. You should create non-*root* filesystems after NSC is installed and configured.



5

Installing NSC Software



Lesson 5

Installing NSC Software

Introduction

Once UnixWare 7.1 is installed in the prescribed manner on the initial root node (node 1) of a cluster, the Non-Stop Cluster (NSC) software is then added to that node. This process *clusterizes* node 1. That is, it converts node 1 from a stand-alone computer running UnixWare 7.1 to the first node in a cluster. Subsequent nodes can then be installed using a NSC-specific simplified procedure.

This lesson shows you how to install the NSC on the initial root node so that it is successfully clusterized and describes how other nodes are added to the cluster.

Objectives

By the end of this lesson, you will be able to:

- Perform NSC pre-installation preparation tasks
- Install NSC software and configure the initial root node
- Test the newly clusterized environment on node 1

AV 5-1

The slide is titled "Installing NSC Software on Node 1" and contains a bulleted list of topics. The footer includes the identifier "Nsc101" and the copyright notice "© The Santa Cruz Operation, Inc.".

Installing NSC Software on Node 1

- Installation Overview
- NSC Installation Issues for Node 1
- Installation Lab

Nsc101 © The Santa Cruz Operation, Inc.

Topics

In order to meet these objectives, this lesson discusses:

- Node 1 Installation Procedure Overview — understanding the steps needed to install NSC on node 1, after the cluster hardware is all assembled and UnixWare 7 has been properly installed on node 1
- NSC Software Installation Information — navigating through the preparation steps and NSC installation prompts for all protected storage methods.
- Lab — installing NSC and related software to clusterize node 1

5.1 Understanding NSC Software Installation

Recall from the overview in Lesson 4: “Installing SCO UnixWare 7,” that your cluster software is assembled and installed in this sequence:

1. Assemble operating system software. Create or obtain vendor-supplied EFS or HBA diskettes, if needed.
2. Install UnixWare 7.
3. Install vendor-supplied EFSes.
4. Prepare the initial root node for NSC installation
5. Install NSC on the initial root node
6. Install the failover root node
7. Install depending nodes, if present.

Just as with the UnixWare 7 installation process, successful installation of the NSC software requires careful preparation. Because installation of all dependent nodes depends on the success of the installation and preparation of the initial root node, you must understand the configuration choices you will make during the NSC installation process before you begin.

This lesson describes Steps 4 through 6 of the software installation procedure. Each section gives you guidance about the selections you need to make during the NSC software preparation and installation steps.

The labs at the end of the lesson provide detailed procedures for installing the initial root node. There are separate procedures within the lab for CNM and shared-storage clusters.

The next lesson instructs you to install the failover root node and includes a lab that demonstrates the failover capability of your newly installed cluster.

5.2 Preparing to Install NSC on the Initial Root Node

Before installing the SCO UnixWare 7 NonStop Cluster software onto the UnixWare 7 initial root node you must perform several preparation tasks in a specific order prior to installing NSC.

This section provides information about the NSC-related pre-installation tasks you need to perform as part of the overall installation procedure included in Section 5.6, “Lab: Installing NSC on the Initial Root Node,” on page 5-18.

5.2.1 Applying Program Technical Fixes (PTFs)

Depending on the release of UnixWare and NSC you are installing, you might need to apply fixes to the operating system *before* installing NSC. For example, the NSC 7.1.1 CD-ROM includes a PTF for UnixWare 7.1.1 that must be installed before you install NSC. This is a package called UW7PTF and is installed using **pkgadd**.

Read the ***NonStop Clusters Release Notes*** for information about PTFs for the releases of UnixWare and NSC used in your classroom.

Note that there might also be fixes available from SCO that are more recent than those included on your NSC CD-ROM. For the latest information specific to your software version, see the SCO Support web site at URL:

<http://www.sco.com/Support/ftplists/nsclist.htm>

Be very careful to distinguish which fixes or supplements are required *before*, and which *after* NSC is installed.

5.2.2 Encapsulating the *root* filesystem

After applying any applicable supplements or PTFs for UnixWare, prepare for NSC installation by encapsulating the *root* filesystem on the internal drive of the initial root node.

The *root* filesystem failover configuration relies on the disk mirroring feature of the Veritas Volume Manager (**VxVM**) feature. Encapsulating the *root* filesystem brings it under Veritas control.

Section 5.6, “Lab: Installing NSC on the Initial Root Node,” on page 5-18 has detailed instructions for using `vxinstall` to encapsulate the *root* filesystem.

Do not install or encapsulate any other drives at this time. If you want to create additional filesystems on the root drive, or place other drives under **VxVM** management, do so after you have installed the NSC software. Adding non-*root* filesystems is discussed in the **UnixWare® 7 NonStop Clusters® Administration** course.

Note After you encapsulate the *root* filesystem, the system *must* be rebooted. Be sure to respond with **y** when prompted to reboot.

5.2.3 Preparing shared storage

When using shared storage units with SCO NonStop Clusters, you might be required to prepare the shared storage before you begin the NSC installation process.

These tasks do not apply to CNM clusters. If your classroom does not use shared storage units, skip this section and proceed to Section 5.3.1, “Selecting optional packages,” on page 5-12.

5.2.3.1 Stamping logical disks on shared storage units

Some shared storage units require you to stamp logical volumes with a unique identifier so they can be recognized by UNIX. This identifier is called a *nexus*.

The nexus is the name of the device in the `/dev/dsk` and `/dev/rdsk` directories and has the form `c#b#t#d#` where:

- The # after the ‘c’ is a controller number
- The # after the ‘b’ is a bus number
- The # after the ‘t’ is a target-id (sometimes called TID) number
- The # after the ‘d’ is a logical unit (sometimes called LUN) number

This must be done after UnixWare 7 is installed on the initial root node, but before the NSC software is installed.

If your cluster uses the Compaq FibreChannel Storage System, you must verify the array’s hardware connections and stamp its logical volumes with a SCSI nexus. Before beginning the NSC installation procedure, read Section 1.9.3, “Stamping disks in a Compaq Fibre array,” on page 65 in the ***NSC Curriculum Hardware Supplement*** for hardware set up information.

If your cluster uses a shared storage unit supplied by another vendor, check the Section 1.9, “Initializing Shared Storage,” in your vendor’s chapter in the ***NSC Curriculum Hardware Supplement*** to see if this requirement applies to your hardware.

You are referred to this section again at the appropriate step in the installation lab in Section 5.6.3, “Procedure for shared-storage clusters,” on page 5-27

5.2.3.2 Creating a UNIX partition and VTOC on shared storage

Some shared storage units require that you create a UNIX partition with a volume table of contents (VTOC) before installing NSC. Otherwise, the NSC installation will fail when trying to move the *root* filesystem from the internal drive to the external storage unit.

Use **diskadd(1M)** to create the UNIX partition with no (0) slices. This creates a valid VTOC table for the drive. The device name argument is of the form *c#b##t##d#*. For example, on a node with:

- An SCSI controller (with the internal root disk) as controller 0
- An IDE controller with a CD-ROM as controller 1
- A Fibre channel HBA connecting to the shared storage unit

This is an example command line for creating a UNIX partition on the first disk on the third controller on node 1:

```
# diskadd c2b0t0d0
```

Use the entire disk for UNIX, unless there is already a system, or vendor, partition on the disk. Do not remove these types of partitions. Use the remainder of the disk for UNIX.

Note Be sure specify 0 as the number of slices on the disk.

5.2.3.3 Controlling shared storage

Some shared storage solutions require that only one node actively controls the shared storage at any time. Since installation of the NSC software includes migrating the *root* filesystem from the internal disk (the one on which you installed UnixWare) to the external shared storage, node 1 must have control of the shared storage for the NSC installation to succeed.

If you are using a shared storage solution where all nodes have active control of the shared storage, such as JBOD or certain Fibre solutions, this is not an issue and you can skip this procedure.

To ensure that node 1 controls the external storage unit, you must power up the cluster hardware in a particular sequence. This is detailed in Section 5.6, “Lab: Installing NSC on the Initial Root Node,” on page 5-18, and summarized here:

1. Turn off all nodes sharing the storage other than node 1.
2. Shutdown node 1 with a normal, full system shutdown, not a shutdown with reboot.
3. Power off node 1.
4. Power off the shared storage devices.
5. Power on the shared storage devices. If the storage boxes and storage controllers are separate, follow the vendor’s instructions for power-on order.
6. Turn on and boot node 1. During the boot, node 1 will gain control of the shared storage.

5.2.4 Removing SDI ghost devices

SDI ghost devices can be created in certain circumstances on UnixWare systems. This occurs in multipath I/O environments—where two host bus adapters in a systems read and write to the same disk

You must remove any ghost devices that exist on your initial root node before installing NSC.

Section 5.6, “Lab: Installing NSC on the Initial Root Node,” on page 5-18 includes the steps used to check for, and remove, ghost devices.

5.3 Installing NSC on the Initial Root Node

Once all preparation tasks have been successfully accomplished, install the NSC software from the NSC CD-ROM using **pkgadd**.

This section provides an overview of all NSC-related information you need to follow the installation procedure provided in Section 5.6, “Lab: Installing NSC on the Initial Root Node,” on page 5-18.

5.3.1 Selecting optional packages

After you start the NSC installation procedure, you might be prompted to select optional packages. This depends on the version of NSC you are installing. Select all optional packages, unless directed otherwise by your instructor.

Read the ***UnixWare NonStop Clusters Release Notes*** for your version of NSC before you make selections to be sure you understand any restrictions or limitations.

5.3.2 Configuring *root* filesystem failover

One of the key features of NSC is the ability to have the cluster continue operation after losing the node controlling the *root* filesystem. This is possible only when there is another node that can take over control of the *root* filesystem.

How you configure failover depends on the type of protected storage method you are using. Regardless of whether you have a CNM or shared-storage cluster, you *must* configure *root* filesystem failover during the NSC installation process.

On CNM clusters:

- If node 1 has only one disk, the NSC installation program assumes you are performing a CNM install.
- For simplicity, do not configure OSU as part of the classroom lab later in this lesson. However, on live systems, it is recommended that you configure OSU.
- If node 1 has multiple disks, select the disk containing the *root* filesystem. For example, if the first disk on the first controller is used for the *root* filesystem, select “Drive 1 - c0b0t0d0” as the device.
- Enter **2** when prompted for the node number of the root failover node.
- Enter **0** when prompted to enter another cluster node number for root failover. This indicates there are no other nodes for root failover.
- The actual migration of the *root* filesystem to node 2 occurs when you install NSC on node 2.

On shared-storage clusters:

- Select the drive number that corresponds to the first external drive in the shared storage unit.
- For simplicity, do not configure OSU as part of the classroom lab later in this lesson. However, on live systems, it is recommended that you configure OSU.
- Enter the node number of the other node connected to the shared storage unit. In your classroom, this is generally node 2.
- Specify whether your shared storage unit is a RAID device.
- For simplicity, do not add additional disks for the *root* filesystem even if your classroom storage unit has multiple disks.
- The *root* filesystem is moved to the shared storage device when you complete this procedure and confirm your selections. This might take a while, depending on the size of your *root* filesystem.

5.3.3 Configuring the CVIP and PIP

During NSC installation, you are prompted configure the *Cluster Virtual IP (CVIP)* for the cluster. The NSC installation procedure automatically promotes the name and IP address of the network interface currently configured on this node to the name and IP address of the cluster.

Then you enter a new host name and IP address for the physical interface (NIC) on this node.

Be prepared with:

- A new host name for this node's network interface
- A new IP address for this node's network interface

The IP address should be on the same subnet as the CVIP.

5.3.4 Configuring optional packages

After you finish configuring CVIP information, the NSC installation program prompts you for any information needed to configure the optional packages you selected to install. The installation program does not actually install the packages until after it installs the NonStop Clusters software.

Follow any instructions that appear during optional package installation. If you are installing the CIM package, you are prompted for this Simple Network Management Protocol (SNMP) information. Enter appropriate values, or press to accept defaults for all fields.

For your classroom configuration, press through each field to accept default values.

5.4 Post-Installation Tasks

Before proceeding to install the dependent nodes in the cluster:

- Install application software you plan to have running on the root node. See the SCO NonStop Cluster Concepts course for help planning your cluster configuration to run applications.

In this course, it is unlikely you will install applications before continuing to Lesson 6: “Installing the Failover Root Node.”

- Remove and reinstall the ODM license using the **License Manager** so the license will apply cluster-wide.

5.5 Lesson Summary

This lesson has discussed how to prepare your initial root node for installing NSC software. An overview of the NSC installation process was provided to familiarize you with the types of configuration questions you need to be prepared to answer, whether your cluster uses CNM or shared storage units for protected storage.

The dependent node installation process was also presented.

Review questions

1. Which filesystems should you encapsulate before installing NSC?
2. How do you configure a NIC on a dependent node?
3. Is split-brain avoidance cabling needed for clusters using CNM?

5.6 Lab: Installing NSC on the Initial Root Node

This lab demonstrates how to install NonStop Cluster software on the initial root node and then test the configuration. Students are divided into two or more groups, one for each cluster.

5.6.1 Overview

Install your initial root node using the steps in the lab applicable to your cluster configuration. Properties of the cluster are established in Lesson 3: “Assembling and Initializing Cluster Hardware.” Review the properties of the cluster including its protected storage method, number of nodes, and SAN technology.

Before you begin, read the **SCO NonStop Clusters Release Notes**, if you have not already done so. You need the latest information applicable to your classroom release of NSC to install the correct PTFs and SLses.

5.6.2 Procedure for CNM clusters

To install NSC software and configure the initial root node:

- Step 1. Install any necessary updates to the operating system provided by SCO.

For example, if you are using UnixWare 7.1 or 7.1.1, install the PTF called UW7PTF from the NSC CD-ROM using **pkgadd**:

pkgadd -d cdrom1 UW7PTF

- Step 2. Configure the LAN NIC for node 1 using the **netcfg(1M)** command. Assign it the host name of the cluster and the IP address you are using as the CVIP.

- Step 3. Encapsulate the *root* filesystem.

- a. As **root**, enter:

vxinstall

A list of disk controllers found on the node is displayed with a description of the disk device-node naming conventions used by **VxVM**.

- b. Press to continue. An information screen is displayed describing the “Quick” and “Custom” options. Press to continue.
- c. Enter **2** to select a “Custom Installation.”
- d. Enter **y** when prompted to encapsulate the boot disk.

This encapsulates all filesystems on the boot disk, including the *root* filesystem. These filesystems will then be VxFS volumes when the node is rebooted.

- e. Accept the default name, *rootdisk*. Press .
- f. Verify the configuration displayed by **vxinstall**. Press to continue.
- g. Press and then enter **4** to leave the other disks unencapsulated.
- h. Press and then enter **4** if **vxinstall** detects other controllers. Leave all other disks on all controllers unencapsulated.

- i. Enter **y** when **vxinstall** displays a list of disks that will be encapsulated. It should only include *c0b0t0d0*.

If it does not, enter **q** to quit and restart **vxinstall**.

- j. Enter **y** when **vxinstall** prompts you to reboot.

WARNING! You *must* enter **y** to reboot after encapsulating the *root* filesystem. Continuing without rebooting the node causes the NSC installation to fail.

The system reboots multiple times. The number of reboots depends upon your system.

- Step 4. When the initial root node finishes booting, switch to the command-line login prompt by pressing the **CTRL** **ALT** **h** key combination, and log in as **root**.

Step 5. Check for SDI ghost devices. Enter:

/etc/scsi/sdighost -l

If this message is displayed, **sdighost** has found no ghost devices and you are finished with this procedure.

UX:sdighost:INFO:Currently, no names are in use as Ghost Names

Step 6. Remove SDI ghost devices, if found:

a. Enter:

/etc/scsi/sdighost -r

b. Shutdown the initial root node. Enter:

/etc/shutdown -i6 -g0 -y

c. Cycle power on the initial root node.

Step 7. When the initial root node finishes booting, switch to the command-line login prompt by pressing the **CTRL ALT h** key combination, and log in as **root**.

Step 8. Start the SCO NSC installation. Insert the NSC CD-ROM in the first CD-ROM drive on node 1 and enter:

pkgadd -d cdrom1 -q NSC

Step 9. Select all optional packages.

Step 10. Identify the node being installed. Since this procedure is always performed on the first node, enter **1** **ENTER**.

Step 11. Select the failover root node. For CNM:

- a. Enter the drive number of the disk that corresponds to the primary disk on the failover root node.
- b. Enter **n** to configure OSU, unless otherwise instructed.
- c. Enter **2** when prompted for the node number of the root failover node.
- d. Enter **0** when prompted to enter another cluster node number for root failover. This indicates there are no other nodes for root failover.

The actual migration of the root filesystem to node 2 will occur when you install NSC on node 2.

Step 12. Configure the CVIP.

- a. Enter **y** when prompted to configure the CVIP. The host name and IP address that was previously configured for the first NIC on node 1 is displayed. This will become the cluster name and CVIP.
- b. Enter the new host name for the NIC.
- c. Enter the new IP address for the NIC.
- d. Verify the values for the new interface to the NIC. Enter **y** if they are correct.

When you confirm the new network information for node 1, a message indicates that the information will be applied when the system is ready to run the NonStop Clusters software.

- Step 13. Configure SNMP information. Press  to accept defaults and continue.
- Step 14. Monitor the messages displayed by the installation program that provide status information about the tasks it is performing, such as copying NonStop Clusters files, mirroring the root filesystem, updating system files, and so on.
- Step 15. Remove the NSC CD-ROM and reboot the node. Enter:
- shutdown -i6 -y -g0**

AV 5-2

Cluster Commands

```
# cluster -v
1: UP
2: UP

# ifconfig -a
lo0: flags=4049<UP,LOOPBACK,RUNNING,MULTICAST> mtu 16384
    inet 127.0.0.1 netmask ffffffff
    inet/perf: rcv size: 4096; send size: 8192; full-size frames: 1
    inet/options: rfc1323
cvip0: flags=84041<UP,RUNNING,MULTICAST,CVIP> mtu 16384
    inet 132.147.119.100 netmask ffffffff
    inet/perf: rcv size: 24576; send size: 24576; full-size frames: 1
    inet/options: rfc1323
net1.2: flags=4043<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 132.147.119.102 netmask ffffffff broadcast 132.147.119.255
    inet/perf: rcv size: 24576; send size: 24576; full-size frames: 1
    inet/options: rfc1323
    ether 00:80:5f:bb:a9:6d
net0.1: flags=44043<UP,BROADCAST,RUNNING,MULTICAST,REMOTE> mtu 1500
    inet 132.147.119.101 netmask ffffffff broadcast 132.147.119.255
    inet/perf: rcv size: 24576; send size: 24576; full-size frames: 1
    inet/options: rfc1323
    ether 00:80:5f:bb:c0:c6
```

Nsc102

© The Santa Cruz Operation, Inc.

AV 5-3

Cluster Commands (Cont'd)

```
# vxprint
Disk group: rootdg

TY NAME      ASSOC      KSTATE  LENGTH  PLOFFS  STATE  TUTILO
PUTILO
dg rootdg    rootdg      -        -        -        -        -
dm nscrootdisk nlc1b0t0d0s0 -      8373375 -        -        -
dm rootdisk  nlc0b0t0d0s0 -      8321670 -        -        -

pl rootvol-01 -          DISABLED 7839720 -        -        -
sd rootdisk-02 rootvol-01  ENABLED 7839720 0        -        -

v rootvol    root        ENABLED 7839720 -        ACTIVE -
pl nscrootvol-01 rootvol    ENABLED 7839720 -        ACTIVE -
sd nscrootdisk-01 nscrootvol-01 ENABLED 7839720 0        -        -

v standvol   fsngen      ENABLED 96390 -        ACTIVE -
pl standvol-01 standvol    ENABLED 96390 -        ACTIVE -
sd rootdisk-03 standvol-01 ENABLED 96390 0        -        -

v swapvol    swap        ENABLED 224910 -        ACTIVE -
pl swapvol-01 swapvol     ENABLED 224910 -        ACTIVE -

# mount
/ on /dev/root read/write/setuid on Mon Apr 17 01:34:23 2000
/dev/fd on /dev/fd read/write on Mon Apr 17 01:34:23 2000
/proc on /proc read/write on Mon Apr 17 01:34:23 2000
/stand/1 on /dev/stand.1 read/write on Mon Apr 17 01:34:24 2000
/system/processor on /processorfs read/write on Mon Apr 17 01:35:05 2000
/dev/tcp on /dev/tcp read/write on Mon Apr 17 01:35:05 2000
/stand/2 on /dev/stand.2 read/write on Mon Apr 17 01:38:59 2000
/stand/3 on /dev/stand.3 read/write on Mon Apr 17 01:39:27 2000
```

Nsc103

© The Santa Cruz Operation, Inc.

Step 16. Verify that the initial root node is successfully running the clusterized operating system. Although there are no other nodes installed yet, you can now see some of the differences in the NSC operating environment, as shown in AV 5-2 and AV 5-3 on page 5-24.

- a. Log into the cluster as **root**.
- b. Use these commands to view cluster-wide devices and information:

```
cluster -v  
ifconfig -a  
vxprint  
mount
```

Step 17. Install any Support Level Supplements (SLS) or PTFs required for NSC.

- a. Check the **SCO NonStop Clusters Release Notes** for any PTFs included on the NSC CD-ROM installation media.
- b. Check the SCO Support web site for additional SLses or PTFs. See URL:

<http://www.sco.com/Support/ftplists/nsclist.htm>

- c. Follow the instructions for installation in the accompanying documentation.

For example, if you are using UnixWare 7.1 or 7.1.1, install both the NSC SLS **nsc1006** from the CD-ROM, and NSC SLS **nsc1009**, which must be downloaded. These are both installed using **pkgadd**:

Follow the prompts to select and install SLS packages.

Reboot when finished:

```
shutdown -i6 -g0 -y
```

Step 18. Install any application software.

In the classroom setting, it is not likely you will be instructed to install additional software. However, in a real-world environment, you will need to do this before proceeding to install dependent nodes.

You have now completed installation of the initial root node.

5.6.3 Procedure for shared-storage clusters

To install NSC software and configure the initial root node in clusters using external shared storage units:

- Step 1. Install any necessary updates to the operating system provided by SCO.

For example, if you are using UnixWare 7.1 or 7.1.1, install the PTF called UW7PTF from the NSC CD-ROM using **pkgadd**. For example, enter:

pkgadd -d cdrom1 UW7PTF

- Step 2. Configure the LAN NIC for node 1 using the **netcfg(1M)** command. Assign it the host name of the cluster and the IP address you are using as the CVIP.

- Step 3. Encapsulate the *root* filesystem.

- a. As **root**, enter:

vxinstall

A list of disk controllers found on the node is displayed with a description of the disk device-node naming conventions used by **VxVM**.

- b. Press to continue. An information screen is displayed describing the “Quick” and “Custom” options. Press to continue.
- c. Enter **2** to select a “Custom Installation.”
- d. Enter **2** to select a “Custom Installation.”
- e. Enter **y** when prompted to encapsulate the boot disk.

This encapsulates all filesystems on the boot disk, including the *root* filesystem. These filesystems will then be VxFS volumes when the node is rebooted.

- f. Accept the default name, *rootdisk*. Press .
- g. Verify the configuration displayed by **vxinstall**. Press to continue.
- h. Press and then enter **4** to leave the other disks unencapsulated.

- i. Press **ENTER** and then enter **4** if **vxinstall** detects other controllers. Leave all other disks on all controllers unencapsulated.
- j. Enter **y** when **vxinstall** displays a list of disks that will be encapsulated. It should only include *c0b0t0d0*.

If it does not, enter **q** to quit and restart **vxinstall**.
- k. Enter **y** when **vxinstall** prompts you to reboot.

WARNING! You *must* enter **y** to reboot after encapsulating the *root* filesystem. Continuing without rebooting the node causes the NSC installation to fail.

The system reboots multiple times. The number of reboots depends upon your system.

- Step 4. When the initial root node finishes booting, switch to the command-line login prompt by pressing the **CTRL** **ALT** **h** key combination, and log in as **root**.
- Step 5. Stamp the logical disks with a SCSI *nexus* if you have a Fibre shared storage unit that does not have SCSI device addresses. See Section 1.9.2, “Stamping disks in the external disk subsystem,” in your vendor’s chapter in the **NSC Curriculum Hardware Supplement** for the procedure specific to your classroom hardware.
- Step 6. Initialize the logical disk to create a VTOC using **diskadd(1M)**. Create a UNIX partition with no slices. For example:

diskadd c2b0t0d0

- a. Enter **y****ENTER** to continue when warned about destroying data.
- b. Enter **y****ENTER** to accept the recommended partitioning of 100% for UNIX, unless your disk already has a vendor or system partition. In that case, allocate the rest of the disk to UNIX.
- c. Enter **n****ENTER** to skip surface analysis.
- d. Enter **0****ENTER** when prompted for the number of slices.
- e. Enter **y****ENTER** to confirm the disk allocation information.

- Step 7. Verify that node 1 has control of the shared storage devices:
- a. Turn off all nodes sharing the storage other than node 1.
 - b. Shutdown node 1 with a normal, full system shutdown, not a shutdown with reboot. For example:

cd /
shutdown -y -g0 -i0
 - c. Power off node 1.
 - d. Power off the shared storage devices.
 - e. Power on the shared storage devices. If the storage boxes and storage controllers are separate, follow the vendor's instructions for power-on order.
 - f. Turn on and boot node 1. During the boot, node 1 will gain control of the shared storage.
 - g. Switch to the home console screen using the   **h** key sequence.
 - h. Log in as **root**.

- Step 8. Check for SDI ghost devices. Enter:

/etc/scsi/sdighost -l

If this message is displayed, **sdighost** has found no ghost devices and you are finished with this procedure.

UX:sdighost:INFO:Currently, no names are in use as Ghost Names

- Step 9. Remove SDI ghost devices, if found:

- a. Enter:

/etc/scsi/sdighost -r

- b. Shutdown the initial root node. Enter:

/etc/shutdown -i6 -g0 -y

- c. Cycle power on the initial root node.

- Step 10. When the initial root node finishes booting, switch to the command-line login prompt by pressing the **CTRL ALT h** key combination, and log in as **root**.

- Step 11. Start the SCO NSC installation. Insert the NSC CD-ROM in the first CD-ROM drive on node 1 and enter:

pkgadd -d cdrom1 -q NSC

- Step 12. Select all optional packages.

- Step 13. Identify the node being installed. Since this procedure is always performed on the first node, enter **1** **ENTER**.

- Step 14. Select the failover root node. For shared storage clusters:
- a. Enter the drive number of the disk that corresponds to the shared storage device. If using RAID, this is a logical disk.
 - b. Enter **n** to configure OSU, unless otherwise instructed.
 - c. Indicate whether the selected disk is a hardware RAID device.
 - d. Enter **n** when prompted to configure additional disks for the *root* filesystem.
 - e. Indicate whether the configuration is correct when the current configuration is summarized.
 - f. Enter **2** when prompted for the node number of the root failover node.
 - g. Enter **0** when prompted to enter another cluster node number for root failover. This indicates there are no other nodes for root failover.
 - h. Wait while the *root* filesystem migrates to the shared storage unit. This takes several minutes.
- Step 15. Configure the CVIP.
- a. Enter **y** when prompted to configure the CVIP. The host name and IP address that was previously configured for the first NIC on node 1 is displayed. This will become the cluster name and CVIP.
 - b. Enter the new host name for the NIC.
 - c. Enter the new IP address for the NIC.
 - d. Verify the values for the new interface to the NIC. Enter **y** if they are correct.
- When you confirm the new network information for node 1, a message indicates that the information will be applied when the system is ready to run the NonStop Clusters software.
- Step 16. Configure SNMP information. Press  to accept defaults and continue.
- Step 17. Monitor these messages displayed by the installation program that provide status information about the tasks it is performing, such as copying NonStop Clusters files, mirroring the *root* filesystem, updating system files, and so on.
-

- Step 18. Remove the NSC CD-ROM and reboot the node. Enter:

shutdown -i6 -y -g0

- Step 19. Verify that the initial root node is successfully running the clusterized operating system. Although there are no other nodes installed yet, you can now see some of the differences in the NSC operating environment, as shown in AV 5-2 and AV 5-3 on page 5-24.

- a. Log into the cluster as **root**.
- b. Use these commands to view cluster-wide devices and information:

cluster -v
ifconfig -a
vxprint
mount

- Step 20. Install any Support Level Supplements (SLS) or PTFs required for NSC.

- a. Check the **SCO NonStop Clusters Release Notes** for any PTFs included on the NSC CD-ROM installation media.
- b. Check the SCO Support web site for additional SLses or PTFs. See URL:

<http://www.sco.com/Support/ftplists/nsclist.htm>

- c. Follow the instructions for installation in the accompanying documentation.

For example, if you are using UnixWare 7.1 or 7.1.1, install both the NSC SLS **nsc1006** from the CD-ROM, and NSC SLS **nsc1009**, which must be downloaded. These are both installed using **pkgadd**:

Follow the prompts to select and install SLS packages.

Reboot when finished:

shutdown -i6 -g0 -y

Step 21. Install any application software.

In the classroom setting, it is not likely you will be instructed to install additional software. However, in a real-world environment, you will need to do this before proceeding to install dependent nodes.

You have now completed installation of the initial root node.

5.7 Relevant Commands and Files

The following commands and files were presented in this lesson.

Relevant commands

cluster (1M)	Displays information about the cluster
fdisk (1M)	Manages UNIX disk partitions
format (1M)	Formats diskettes
ifconfig (1M)	Configures network interfaces
mount (1M)	Mounts filesystems
pkgadd (1M)	Installs software
sdghost (1Mtcp)	Displays ghost devices
shutdown (1M)	Changes system/cluster state
vxinstall (1M)	Encapsulates filesystems to put them under Veritas Volume Manager control
vxprint (1M)	Displays information about Veritas volumes and plexes

Relevant files

<i>/tmp</i>	Is the <i>/tmp</i> filesystem
<i>/var/tmp</i>	Is the <i>/var/tmp</i> filesystem
<i>/var</i>	Is the <i>/var</i> filesystem

5.8 Answers to Review Questions

This section contains the answers to exercises, review questions, and labs presented in this lesson.

Review Question Answers

This section contains the answers to the review questions on page 5-17.

1. Encapsulate only the root filesystem before installing NSC. Configure other non-root filesystems you want to put under control of the Veritas Volume Manager after you have successfully installed NSC on the initial root node.
2. Network interfaces are always configured using **netcfg**, whether adding a NIC to a node, or a CVIP to the cluster.
3. Split-brain avoidance should be configured on all CNM clusters because they do not use SAN switches.



6

Installing the Failover Root Node



Lesson 6

Installing the Failover Root Node

Introduction

Once UnixWare 7.1 is installed in the prescribed manner on the initial root node (node 1) of a cluster, the Non-Stop Cluster (NSC) software is then added to that node. This process *clusterizes* node 1. That is, it converts node 1 from a stand-alone computer running UnixWare 7.1 to the first node in a cluster. Subsequent nodes can then be installed using a NSC-specific simplified procedure.


This lesson shows you how to install the NSC on the initial root node so that it is successfully clusterized and describes how other nodes are added to the cluster.

Objectives

By the end of this lesson, you will be able to:

- Perform NSC pre-installation preparation tasks
- Install NSC software and configure the failover root node
- Simulate node failure to demonstrate the NSC failover capability

AV 6-1



Topics

In order to meet these objectives, this lesson discusses:

- Node 2 Installation Procedure Overview — understanding the steps needed to install NSC on node 2, the failover root node
- NSC Dependent Node Installation Information — navigating through the preparation steps and NSC installation prompts
- Lab — installing NSC and related software on the failover root node and simulating node failure

6.1 Understanding NSC Dependent Node Installation

Recall from previous lessons that your cluster software is assembled and installed in this sequence:

1. Assemble operating system software. Create or obtain vendor-supplied EFS or HBA diskettes, if needed.
2. Install UnixWare 7.
3. Install vendor-supplied EFSes.
4. Prepare the initial root node for NSC installation
5. Install NSC on the initial root node
6. Install the failover root node
7. Install other depending nodes, if present.

After your initial root node is completely installed and you have performed all post-installation tasks described in Lesson 5: “Installing NSC Software,” you can install dependent nodes. This is the last step in the installation sequence.

The failover root node is always installed immediately after the initial root node. The installation process for all dependent nodes is simple and fast. Most of the configuration decisions affecting the installation of dependent nodes was made during the installation of the initial root node. There are three basic steps:

1. Create the dependent node boot diskettes (DNBD).
2. Boot from the floppies and respond to a series of prompts.
3. Reboot the node.

This lesson describes the dependent node installation process in detail. Each section gives you guidance about the selections you need to make during the NSC dependent node software preparation and installation steps.

The lab at the end of the lesson provides detailed procedures for installing the failover root node. There are separate procedures within the lab for CNM and shared-storage clusters.

The final lab demonstrates the failover capability of your newly-installed cluster.

6.2 Installing the Failover Root Node

This section provides an overview of all NSC-related information you need to perform the installation procedure provided in Section 6.4, “Lab: Creating Dependent Node Boot Diskettes,” on page 6-13.

6.2.1 Creating dependent node boot diskettes

When NSC is installed on your initial root node, a program called **make_floppy** is added. This is used to create pairs of dependent node boot diskettes used to install each of the dependent nodes in your cluster.

The **make_floppy** command

- Creates a bootable diskette image containing a minimal NSC kernel
- Creates a diskette image of a minimal NSC *root* filesystem
- Writes the disk images to floppy diskettes.

These disks are analogous to boot and root floppies used to start the general installation procedure for many version of UNIX.

Note Before you run **make_floppy**, you must prepare two formatted diskettes. These diskettes must be error-free or **make_floppy** will fail when attempting to write to the diskettes. If it fails, you must start again from the beginning and recreate the images. This is a time-consuming process. To help insure that **make_floppy** succeeds the first time, use the **verify** option to format.

Some versions of **make_floppy** have undocumented flags to store the DNBD images on the hard disk. Then, if the write process fails, you can restart from that point without having to recreate the images. View the **make_floppy** shell script to see if the version you are using supports these flags. If so, you are recommended to run **make_floppy** to save the images on the hard disk first. Then run **make_floppy** again to write the diskettes.

6.2.2 Booting the failover root node using the DNBDs

The installation of dependent nodes is started by booting the node using DNBD1 diskette. After you insert the DNBD2 diskette, the installation process starts automatically.

The next sections describe the information you need to respond to installation prompts.

6.2.3 Adding the node license

You must add a UnixWare 7 license when installing each dependent node. This should be the same type of license Edition as the root node, but a unique UnixWare 7 license for this node. The **License Manager** is automatically started during the NSC dependent node installation process. Add the license number, data, and node number in the **License Manager** form.

Be sure you enter a unique UnixWare 7 license. Do *not* enter the NSC license.

6.2.4 Checking the fdisk table for CNM

The disk in the failover root node that contains the failover *root* filesystem must have a valid partition table, if you have a two-node cluster using CNM. It must also have enough space for a UNIX partition large enough to hold a mirror of *root* filesystem on node 1.

Hence, this procedure must be done *from node 1*, just after adding the license during the dependent node installation process..

Note The timing for performing this step is critical. It must be performed after node 2 has booted and it's root disk device is available, but before the step in the NSC installation which configures the disk on node 2. Otherwise, the dependent node installation process will fail at this step.

6.2.5 Configuring external storage for dependent nodes

When using an external shared storage unit for your *root* filesystem, the NSC installation process configures the dependent node disks.

Select the option for automatic configuration and apply changes when prompted.

6.2.6 Configuring split-brain avoidance for CNM

Recall that a split-condition can exist when both nodes try to act as active root nodes. Certain SAN implementations that do not use SAN switches can experience this condition. If this occurs, your cluster can experience these problems:

Data corruption	Both the root and failover nodes write to the <i>root</i> filesystem on their local disks and the <i>root</i> filesystem becomes unsynchronized
Ping-pong panics	Each node believes it cannot take ownership of the root filesystem, because their mirror is not up to date, and panics

You can avoid this by configuring an alternate communication path using a null-modem cable between serial ports on two nodes. This prevents your *root* filesystem from being destroyed when the SAN communication is lost.

You are recommended to configure and test split-brain avoidance whenever you are installing a CNM cluster. You must have the serial cable attached to successfully complete the test. However, even if you do not have the cable attached, configure the NSC software for split-brain avoidance and test it later.

6.2.7 Configuring networking on the dependent node

During NSC installation, you are prompted to configure any network adapters present in the dependent node. The installation process automatically starts the **Network Configuration Manager**.

Configure networking for the NICs in the dependent node just as you would in any UnixWare system

Remember these points when configuring this NIC:

- If the dependent node has multiple NICs and they support the standard UnixWare 7 failover feature, you have the option to configure one as a *standby* interface. For NSC, do not configure NICs as standby. All NICs can be active in an NSC environment, and failover is achieved by way of the CVIP facility.
- The IP address for this NIC must be in the same subnet as the CVIP in order to support failover from another node.
- The host name should be unique.

Note The **Network Configuration Manager** automatically enters the cluster virtual host name in the TCP/IP configuration form. Change this to a unique name for this node.

6.2.8 Installing the Mirroring license

For CNM clusters, you must remove and reinstall the Veritas Mirroring or ODM license after the cluster is fully installed.

If the license is not applied cluster-wide, the root filesystem cannot be remirrored in some circumstances. This will cause ping-pong panics during boot when the root node has failed over.

Use the SCOadmin **License Manager** to remove and reinstalled the Veritas Mirroring or ODM license.

6.3 Lesson Summary

This lesson has discussed how to prepare your initial root node for installing NSC software. An overview of the NSC installation process was provided to familiarize you with the types of configuration questions you need to be prepared to answer, whether your cluster uses CNM or shared storage units for protected storage.

The dependent node installation process was also presented.

Review questions

1. Which filesystems should you encapsulate before installing NSC?
2. How do you configure a NIC on a dependent node?
3. Is split-brain avoidance cabling needed for clusters using CNM?

6.4 Lab: Creating Dependent Node Boot Diskettes

When the initial root node installation is finished, you are ready to install the cluster's failover root node. The installation procedure for dependent nodes is much less complex than the setup of node 1.

6.4.1 Overview

Follow this procedure to create the DNBDs needed to install each dependent node.

If your classroom cluster has more than two nodes, make a separate set for each additional dependent node.

After creating the boot diskettes, install node 2 as described in Section 6.5, "Lab: Installing the Failover Root Node," on page 6-15.

6.4.2 Procedure: Using `make_floppy` to create DNBDs

NonStop Clusters software includes a **make_floppy** command to make bootable versions of the single-system image containing all the information you provided in previous installation lessons. You then use these diskettes to install each dependent node in your cluster.

To make the dependent node boot diskettes:

- Step 1. Format two 1.44MB floppies using the **format(1M)** command. Although it is slow, use the verify option to insure that the floppies are error-free. Enter:

```
format -v /dev/rdisk/f03ht
```

- Step 2. Label the floppies *DNBD1* and *DNBD2*.

- Step 3. On node 1, insert DNBD1 into the diskette drive.

- Step 4. Run **make_floppy**. Enter:

```
make_floppy diskette1
```

Wait while the boot image is built. This will take several minutes.

- Step 5. When prompted, replace DNBD1 with DNBD2 and press .

- Step 6. When the **make_floppy** command completes, remove DNBD2 from the drive.

- Step 7. Write protect both diskettes.

When finished, install the failover root node according to the procedure in Lesson 6: “Lab: Installing the Failover Root Node.”

6.5 Lab: Installing the Failover Root Node

After creating the dependent node boot diskettes, you are ready to install the cluster's failover root node. The installation procedure for dependent nodes is much less complex than the setup of node 1.

Separate procedures are provided for CNM and shared storage clusters.

6.5.1 Overview

This lab has separate procedures for each type of cluster, CNM or shared storage. Use the lab applicable to your classroom configuration.

After installing and testing dependent nodes, verify that root failover is working by completing the final lab, Section 6.6, "Lab: Testing Failover," on page 6-23.

6.5.2 Procedure: Installing the failover root node for CNM clusters

To install the failover root node of a CNM cluster:

- Step 1. Boot the initial root node, if it is not already running.
- Step 2. Initialize the failover root node hardware, if you have not already done so. See Section 1.8, “Initializing Nodes,” in your vendor’s chapter in the *NSC Curriculum Hardware Supplement*.
- Step 3. Boot the dependent node using the first DNBD diskette.
- Step 4. Insert the second DNBD when prompted and press **ENTER**.
- Step 5. Enter node number **2** for the failover root node, when prompted.
- Step 6. Ensure that the disk on this node has a valid **fdisk** table.
- Log in as **root** on node 1.
 - Change to the `/dev/dsk` directory.
 - View a listing of all disk devices on node 2. Enter:

ls n2*
 - Note the device node that has the lowest number after the letter c and the rest of the string containing `b0t0d0s0`. For example, `n2c0b0t0d0s0` or `n2c1b0t0d0s0`.
 - Use **fdisk**(1M) with the device you selected to see whether a UNIX partition exists on that device. For example:

fdisk n2c0b0t0d

If the disk has a UNIX partition, skip the rest of this procedure.
 - Use **diskadd**(1M) with the device you selected to create a UNIX partition with 0 slices on that device if a partition does not exist. For example:

diskadd n2c0b0t0d0

Notice that the `s0` portion is left off.

Be careful not to delete any system partition that may have been created by setup software proprietary to the vendor of the node. Other than that, the rest of the disk should be used for a single UNIX partition.

Step 7. Enter an appropriate UnixWare 7 license for this node.

- a. Enter **y** to start the **License Manager** and press **F1** or tab to access the menu bar.
- b. Select **License** ⇨ **Add**.
- c. Enter this node's UnixWare license number, code and data from the COLA that came in your *SCO UnixWare* media kit.

Recall, the dependent nodes must have the same *type* of license Edition as the initial root node, but the actual license number must be unique. If you are not sure which license is required, see Section 4.3.1, "UnixWare 7 licenses," on page 4-8.

Enter the node number for this dependent node.

- d. Tab to the **License** button and press **ENTER**.
- e. Press **ENTER** to select the **OK** button when prompted.

Notice the **License Manager** is clusterized, and shows all licenses for each node.

- f. Select **Host** ⇨ **Exit** to quit.

- Step 8. Configure split-brain avoidance, if using CNM:
- a. Enter **y** when prompted to configure split-brain avoidance.
 - b. Select the port on node 1 to which the serial cable is attached.
 - c. Select the port on node 2 to which the serial cable is attached.
 - d. Enter **y** when prompted to confirm your choices.
 - e. Attach the cable to the selected ports, if you have not already done so during hardware assembly.
 - f. Test the connection.
- Step 9. Configure the LAN NIC cards on the dependent node.
- a. Enter **y** when prompted to enter network information.

The **Network Configuration Manager** is started. It displays a list of all network interface controllers (NICs) and CVIP(s) already installed in the cluster.
 - b. Press  (tab) ⇨  to access the **Hardware** menu item.
 - c. Select **Add new LAN adapter**. The **Add new LAN adapter** window appears. This window contains the list of LAN adapters not yet configured.

Do not configure any NICs as *standby* if your node has multiple NICs.
 - d. Use your arrow keys to highlight the LAN NIC are configuring.
 - e. Press the  (tab) twice to highlight **Continue** and press . Wait while the network interface is added to the cluster.
 - f. If the driver for the NIC has special configuration capabilities, another window appears asking you to configure those capabilities. Configure the capabilities now, if needed.
 - g. The **Add Protocol** window appears with the **TCP/IP** protocol already selected (if it is not, select it).
 - h. Press  (tab) to select **Add**. This displays the **Internet Protocol Configuration** window.

- i. Enter the appropriate information for the node you are installing.

Note Be sure to enter a new host name. By default, the cluster name is entered in the “Host name” field. Change this to a node-specific name.

Choose an IP address on the same subnet as the CVIP and PIPs in the initial root node. Otherwise, the CVIP will not failover if the root node fails.

- j. Press  (tab) to select **OK**. A window displays a successful completion message.
- k. Press  to select **OK**. You are returned to the main **netcfg** window.
- l. Configure additional NICs, if this node has multiple network adapters.
- m. From the **Hardware** menu item, press  and select **Exit**.

Step 10. Remove any diskettes from the diskettes or CD-ROMs and then press  to reboot the cluster.

Step 11. Test the cluster. Log into the cluster as **root** and use these commands to view cluster-wide devices and information:

```
cluster -v  
ifconfig -a  
vxprint  
mount
```

Step 12. Run any vendor-supplied cluster test utilities. See Section 1.12, “Vendor-Supplied Software,” in your vendor’s chapter in the *NSC Curriculum Hardware Supplement*.

Step 13. Use the SCOadmin **License Manager** to remove and reinstalled the Veritas Mirroring or ODM license.

6.5.3 Procedure: Installing the failover root node for shared-storage clusters

To install the failover root node for shared storage clusters:

- Step 1. Boot the initial root node, if it is not already running.
- Step 2. Initialize the failover root node hardware, if you have not already done so. See Section 1.8, “Initializing Nodes,” in your vendor’s chapter in the *NSC Curriculum Hardware Supplement*.
- Step 3. Boot the dependent node using the first DNBD diskette.
- Step 4. Insert the second DNBD when prompted and press .
- Step 5. Enter node number **2** for the failover root node, when prompted.
- Step 6. Enter an appropriate UnixWare 7 license for this node.
 - a. Enter **y** to start the **License Manager** and press or tab to access the menu bar.
 - b. Select **License** ⇒ **Add**.
 - c. Enter this node’s UnixWare license number, code and data from the COLA that came in your *SCO UnixWare* media kit.

Recall, the dependent nodes must have the same *type* of license Edition as the initial root node, but the actual license number must be unique. If you are not sure which license is required, see Section 4.3.1, “UnixWare 7 licenses,” on page 4-8.

Enter the node number for this dependent node.
 - d. Tab to the **License** button and press .
 - e. Press to select the **OK** button when prompted.

Notice the **License Manager** is clusterized, and shows all licenses for each node.
 - f. Select **Host** ⇒ **Exit** to quit.

- Step 7. Set up the dependent node's disk.
- Select "Disk 1" when prompted.
 - Select "Automatic configuration."
 - Enter the number of the node you are installing, when prompted with the "Node X Disk Setup" screen.
 - Select "Apply changes to hard disk(s)."
 - Select "Apply changes and exit."
- Step 8. Configure the LAN NIC cards on the dependent node.
- Enter **y** when prompted to enter network information.

The **Network Configuration Manager** is started. It displays a list of all network interface controllers (NICs) and CVIP(s) already installed in the cluster.
 - Press  (tab) ⇨  to access the **Hardware** menu item.
 - Select **Add new LAN adapter**. The **Add new LAN adapter** window appears. This window contains the list of LAN adapters not yet configured.

Do not configure any NICs as *standby* if your node has multiple NICs.
 - Use your arrow keys to highlight the LAN NIC are configuring.
 - Press the  (tab) twice to highlight **Continue** and press . Wait while the network interface is added to the cluster.
 - If the driver for the NIC has special configuration capabilities, another window appears asking you to configure those capabilities. Configure the capabilities now, if needed.
 - The **Add Protocol** window appears with the **TCP/IP** protocol already selected (if it is not, select it).
 - Press  (tab) to select **Add**. This displays the **Internet Protocol Configuration** window.

- i. Enter the appropriate information for the node you are installing.

Note Be sure to enter a new host name. By default, the cluster name is entered in the “Host name” field. Change this to a node-specific name.

Choose an IP address on the same subnet as the CVIP and PIPs in the initial root node. Otherwise, the CVIP will not failover if the root node fails.

- j. Press  (tab) to select **OK**. A window displays a successful completion message.
- k. Press  to select **OK**. You are returned to the main **netcfg** window.
- l. Configure additional NICs, if this node has multiple network adapters.
- m. From the **Hardware** menu item, press  and select **Exit**.

Step 9. Remove any diskettes from the diskettes or CD-ROMs and then press  to reboot the cluster.

Step 10. Test the cluster. Log into the cluster as **root** and use these commands to view cluster-wide devices and information:

```
cluster -v  
ifconfig -a  
vxprint  
mount
```

Step 11. Run any vendor-supplied cluster test utilities. See Section 1.12, “Vendor-Supplied Software,” in your vendor’s chapter in the *NSC Curriculum Hardware Supplement*.

6.6 Lab: Testing Failover

This lab simulates a hardware failure that causes node 1 to become unusable and shows how control of the root filesystem migrates to node 2.

6.6.1 Overview

Node 1 will be powered down to simulate a hardware failure. You then log in to node 2 and confirm that the root filesystem is under control of node 2, the failover root node.

WARNING! If your cluster uses CNM, you must ensure that the failover root node filesystem mirror is current before you force a failure on node 1. Otherwise, node 2 will panic during boot.

6.6.2 Procedure: Simulating root node failure

This procedure can be used to simulate a hardware failure of the root node. This demonstrates the failover process where the failover root node takes control of the root filesystem.

- Step 1. Log into the cluster as **root**.
- Step 2. If your cluster uses CNM, verify that the *root* filesystem mirror on the failover root node is up to date using **vxinfo**. Enter:

```
vxinfo -p | grep IOFAIL
```

If the *rootvol-02* plex has **IOFAIL** as the mode, the failover *root* filesystem mirror is not up to date. **DO NOT PROCEED** until the mirror is updated. The **vxrecover** process updates the mirror. Make sure it has completed, and recheck the status using **vxinfo** before you continue.

WARNING! If you power off your root node before the *root* filesystem mirror is synchronized, your failover root node will panic when booted. Make sure you have removed and reinstalled the Mirroring or ODM license before proceeding.

- Step 3. Determine which node controls the *root* filesystem. This node is the root node. Enter:

```
where /
```

- Step 4. Log into the console on node 2.



- Step 5. Make sure the *root* filesystem mirror is up to date before proceeding. This is described in Step 2.

- Step 6. Run a command that will continue for several minutes. For example:

```
find / -name * print.
```

Step 7. Power off the root node. Note what happens on node 2. Your shell and viewer process continue to run, although it might pause during failover.

Step 8. Verify that the root filesystem is now controlled by the failover root node.
Enter:

where /

The new root node (formerly the failover root node) now has control of the *root* filesystem.

6.7 Relevant Commands and Files

The following commands and files were presented in this lesson.

Relevant commands

cluster(1M)	Displays information about the cluster
format(1M)	Formats diskettes
ifconfig(1M)	Configures network interfaces
make_floppy(1M)	Makes boot disks for installing dependent cluster nodes
mount(1M)	Mounts filesystems
pkgadd(1M)	Installs software
shutdown(1M)	Changes system/cluster state
vxinstall(1M)	Encapsulates filesystems to put them under Veritas Volume Manager control
vxprint(1M)	Displays information about Veritas volumes and plexes

6.8 Answers to Review Questions

This section contains the answers to exercises, review questions, and labs presented in this lesson.

Review Question Answers

This section contains the answers to the review questions on page 6-12.

1. Encapsulate only the *root* filesystem before installing NSC. Configure other non-root filesystems you want to put under control of the Veritas Volume Manager after you have successfully installed NSC on the initial root node.
2. Network interfaces are always configured using **netcfg**, whether adding a NIC to a node, or a CVIP to the cluster.
3. Split-brain avoidance should be configured on all CNM clusters because they do not use SAN switches.

