HP NonStop S-Series System Expansion and Reduction Guide

Abstract

This guide describes how to expand or reduce an HP NonStop™ S-series system by adding or removing enclosures.

Product Version

N.A.

Supported Release Version Updates (RVUs)

This publication supports G06.24 and all subsequent G-series RVUs until otherwise indicated by its replacement publication.

Part Number	Published
522465-009	July 2007

Document History

Part Number	Product Version	Published
522465-002	N.A.	May 2002
522465-003	N.A.	August 2002
522465-004	N.A.	May 2003
522465-005	N.A.	September 2003
522465-006	N.A.	September 2004
522465-009	N.A.	July 2007

HP NonStop S-Series System Expansion and Reduction Guide

Glossary | Index | Figures | Tables

7. Troubleshooting

A. Common System Operations

Figures

Tables

What's New in This Guide

Guide Information

Abstract

This guide describes how to expand or reduce an HP NonStop™ S-series system by adding or removing enclosures.

Product Version

N.A.

Supported Release Version Updates (RVUs)

This publication supports G06.24 and all subsequent G-series RVUs until otherwise indicated by its replacement publication.

Part Number	Published
522465-009	July 2007

Document History

Part Number	Product Version	Published
522465-002	N.A.	May 2002
522465-003	N.A.	August 2002
522465-004	N.A.	May 2003
522465-005	N.A.	September 2003
522465-006	N.A.	September 2004
522465-009	N.A.	July 2007

New and Changed Information

These changes have been made to this version of the manual:

Section	Change
Section 4, Reducing a System Online	• Connections between router ports and the ServerNet cable connectors on SEBs and MSEBs have been added to 8.3 Disable ServerNet Ports in Y-fabric MSEBs (or SEBs) on page 4-12 • Connections between router ports and the ServerNet cable connectors on SEBs and MSEBs have been added to 8.6 Disable ServerNet Ports in X-fabric MSEBS on page 4-15
Appendix B, ServerNet Cabling	• A new illustration and tables have been added to show connections to SEBS and MSEBs.

About This Guide

This guide describes how to expand or reduce a NonStop S-series system by adding or removing enclosures.

This guide does not describe operations information for ServerNet clusters. For information regarding the installation, configuration, and management of HP NonStop ServerNet cluster hardware and software, see the *ServerNet Cluster Manual* and the *ServerNet Cluster 6780 Planning and Installation Guide*.

Who Should Read This Guide

This guide is written for customers and service providers preparing to increase or reduce the number of enclosures in a NonStop S-series system.

Section Abstracts

Where to Get More Information

Manuals

For abstracts of the NonStop S-series manuals, see the *NonStop S-Series Planning and Configuration Guide*.

CSSI Web

The CSSI Web provides procedures, part numbers, troubleshooting tips, and tools for servicing NonStop S-series systems.

A link to the CSSI web can be found in the left navigation area of the NTL home page.

Authorized service providers can also order the CSSI Web and HSM CD:

- Channel Partners and Authorized Service Providers: Order the CD from the SDRC at https://scout.nonstop.compaq.com/SDRC/ce.htm.
- HP employees: Subscribe at World on a Workbench (WOW). Subscribers automatically receive CD updates. Access the WOW order form at http://hiimpact.americas.cpqcorp.net/wow/order.asp.

Note. After G06.23, a snapshot of the CSSI Web is no longer included on the system console installer CD.

Guided Replacement Procedures

This guide refers to the guided replacement procedures. These automated tools guide you step by step through replacing many customer-replaceable units (CRUs).

Guided procedures are part of the HP Open Systems Management (OSM) package and the Compaq TSM package. OSM is the system management tool of choice for HP NonStop S-series systems.

OSM Guided Procedures

In OSM, guided procedures are launched by actions within the OSM Service Connection. To access the OSM guided replacement procedures:

1. Log on to the OSM Service Connection.
2. Locate and right click on the object in the Tree pane.
3. Select Actions.
4. Select **Replace** from among the available actions for that resource.
5. Click **Perform action**; the action launches the guided procedure.

See to the *OSM User's Guide* for a list of the OSM guided procedures.

TSM Guided Procedures

To access the TSM guided replacement procedures:

Start > Programs > Compaq TSM > Guided Replacement Tools

Note. The Guided Replacement Toolkit (GRT) is used to replace and an IOMF, PMF, power supply, or 6760 ServerNet Device Adapter (ServerNet/DA) in a system running TSM product version T7945AAW (shipped with the G06.12 RVU) or earlier.

Notation Conventions

Hypertext Links

Blue underline is used to indicate a hypertext link within text. By clicking a passage of text with a blue underline, you are taken to the location described. For example:

This requirement is described under Backup Physical Disk Drives on page 3-2.

General Syntax Notation

The following list summarizes the notation conventions for syntax presentation in this manual.

UPPERCASE LETTERS. Uppercase letters indicate keywords and reserved words; enter these items exactly as shown. Items not enclosed in brackets are required. For example:

```
MAXATTACH
```

lowercase italic letters. Lowercase italic letters indicate variable items that you supply. Items not enclosed in brackets are required. For example:

file-name

computer type. `Computer type` letters within text indicate C and Open System Services (OSS) keywords and reserved words; enter these items exactly as shown. Items not enclosed in brackets are required. For example:

```
myfile.c
```

italic computer type. *`Italic computer type`* letters within text indicate C and Open System Services (OSS) variable items that you supply. Items not enclosed in brackets are required. For example:

pathname

[ ] Brackets. Brackets enclose optional syntax items. For example:

```
TERM [\system-name.]$terminal-name
INT[ERRUPTS]
```

Punctuation. Parentheses, commas, semicolons, and other symbols not previously described must be entered as shown. For example:

```
error := NEXTFILENAME ( file-name ) ;
LISTOPENS SU $process-name.#su-name
```

Item Spacing. Spaces shown between items are required unless one of the items is a punctuation symbol such as a parenthesis or a comma. For example:

```
CALL STEPMOM ( process-id ) ;
```

If there is no space between two items, spaces are not permitted. In the following example, there are no spaces permitted between the period and any other items:

```
$process-name.#su-name
```

Notation for Messages

The following list summarizes the notation conventions for the presentation of displayed messages in this manual.

Nonitalic text. Nonitalic letters, numbers, and punctuation indicate text that is displayed or returned exactly as shown. For example:

```
Backup Up.
```

lowercase italic letters. Lowercase italic letters indicate variable items whose values are displayed or returned. For example:

```
p-register
process-name
```

[ ] Brackets. Brackets enclose items that are sometimes, but not always, displayed. For example:

```
Event number = number [ Subject = first-subject-value ]
```

A group of items enclosed in brackets is a list of all possible items that can be displayed, of which one or none might actually be displayed. The items in the list might be arranged either vertically, with aligned brackets on each side of the list, or horizontally, enclosed in a pair of brackets and separated by vertical lines. For example:

```
proc-name trapped [ in SQL | in SQL file system ]
```

Change Bar Notation

Change bars are used to indicate substantive differences between this edition of the manual and the preceding edition. Change bars are vertical rules placed in the right margin of changed portions of text, figures, tables, examples, and so on. Change bars highlight new or revised information. For example:

> When you expand your system online, you configure enclosures individually and then cable them together into blocks before adding them to your system. You add blocks to the system one block at a time, no matter how large or small each block is.

1 The Resizing Process

This section gives background information necessary to understand system resizing.

Definition

System resizing is the process of doing either of the following to a NonStop S-series system:

- Expanding the system by adding enclosures
- Reducing the system by removing enclosures

Prerequisites

To resize a NonStop S-series system online:

You must be familiar with:	• OSM or TSM software package • HP NonStop Kernel operating system • Subsystem Control Facility (SCF) • NonStop S-series system
The system you are planning to resize must be:	• Fully functional • Currently running the following software and firmware if you are planning to **reduce** the system online: ◦ G06.04 or later RVU ◦ Compatible version of the OSM or TSM package ◦ OSM running on G06.08 or later RVU
The system you are planning to resize must have:	Identical firmware and boot code on all: • Processors • Service processors (SPs) • SCSI controllers This guide describes how to make firmware and boot code identical.

Also, it is recommended, but not required, that you complete a NonStop S-series support class before you resize a system.

Expansion With IOAM Enclosures

You can use I/O Adapter Module (IOAM) enclosures installed in 19-inch racks or modular cabinets to expand your system. An IOAM enclosure connects to the MSEB of S76000 and later S-series systems.

Each IOAM enclosure can accept up to 10 ServerNet adapters designed specifically for the IOAM enclosure. These include the Fibre Channel ServerNet adapter (FCSA) and the Gigabit Ethernet 4-Port ServerNet Adapter (G4SA).

FCSAs provide access to Fibre Channel storage devices, such as an Enterprise Storage System (ESS). For more information, see the *Fibre Channel ServerNet Adapter Installation and Support Guide.*

G4SAs provide increased Ethernet capacity. An IOAM enclosure with 10 G4SAs provides up to 40 ports of Gigabit Ethernet connections. For more information, see the *Gigabit Ethernet 4-Port ServerNet Adapter Installation and Support Guide.*

△ **Caution.** IOAM enclosures and components are FRUs and must be installed by service providers trained by HP.

For planning purposes, an IOAM enclosure is equivalent to an S-series I/O enclosure. In Section 2, Planning System Expansion, examples of blocks using I/O enclosures also apply to IOAM enclosures. However, IOAM enclosures may only be able to occupy certain positions in the topology. See you service provider for details.

IOAM enclosures are physically different from I/O enclosures and contain different components with different slot numbering. IOAM enclosures also require the use of a maintenance switch to communicate with the Maintenance Entity (ME) in each module of an IOAM enclosure.

If you are resizing your system with IOAM enclosures, your service provider needs to consult these publications:

- *Modular I/O Installation and Configuration Guide* (in the NTL Hardware Service and Maintenance collection)
- *NonStop S-Series Planning and Configuration Guide*
- *NonStop S-Series Hardware Installation and FastPath Guide*

Planning for Resizing

Before you resize your system, review your system planning and make new plans if necessary. The issues to consider before you resize your system include (but are not limited to):

Does Your System Have Enough MSEBs and MSEB PICs?

In any system that is expanded, the number of modular ServerNet expansion boards (MSEBs) and MSEB Plug-in Cards (PICs) must be adequate to support the increased number of processors in the system after expansion. Your system probably contains only the correct number of MSEBs and MSEB PICs for its current configuration.

Before you expand your system, be sure you have the appropriate number of MSEBs and MSEB PICs.

For Information About	See
MSEBs	*NonStop S-Series Planning and Configuration Guide*
Installing MSEBs	The Replace SEB or MSEB guided procedure and related procedure help
Adding PICS to an MSEB	CSSI Web

△ **Caution.** To add more PICs to an MSEB, run the Replace SEB or MSEB guided procedure to remove the MSEB. The guided procedure ensures that the MSEB is ready to remove from the system. MSEB PICs are Class-3 CRUs: you must be trained by HP to install them. Be sure you are running T7945G06ABF or later if you plan to add PICS to an MSEB.

Should You Install a New RVU?

You can resize your system online only if your system is running the appropriate RVU:

- For online expansion: G06.00 or later, with a compatible TSM package
- For online reduction: G06.04 to G06.08, with a compatible TSM package
- For online expansion and reduction: G06.08 or later, with a compatible OSM or TSM package
- For expansion involving IOAM enclosures installed in modular racks: G06.24 or later, with a compatible OSM package

△ **Caution.** IOAM enclosures and components must be installed by service providers trained by HP.

If your system is not currently running an appropriate RVU for the resizing actions you plan to take, you must halt your system to install an appropriate RVU. In this case, choose between:

- Installing an appropriate RVU and then resizing the system while the system is stopped (keeps your system offline longer)
- Installing and loading an appropriate RVU, and performing the resizing procedure while the system remains online (maximizes availability)

Topology

The topology of your system is the basic structure of your system.

Your system can be configured for either of two topologies: **Tetra 8** or **Tetra 16**. Each topology has advantages, depending on how you use your system and how you plan to resize it in the future.

Note. It is important to decide ahead of time which system topology you want in your resized system because topologies cannot be changed online. After you configure your system for a particular topology, changing the topology requires stopping your system.

If you are expanding your system, you might need a different topology. For example, if your system is a Tetra 8 topology with the maximum number of I/O and processor enclosures attached, you cannot expand it online. You must reconfigure it offline as a Tetra 16 before it can support more enclosures.

For Information About	See
Topologies	*NonStop S-Series Planning and Configuration Guide*

Choosing a Topology

NonStop S-series servers are shipped with Tetra 8 topology unless the system ordered specifically requires more than eight processors or needs to be able to grow online to more than eight processors.

When planning to resize your system, consider which of the following cases your system fits:

Case	Description	Examples
1	The system's topology does not need to change.	• The system starts as a Tetra 8 topology and will never get too large to be a Tetra 16 topology. • The system starts and stays as a Tetra 16 topology. There is no advantage to changing to a Tetra 8 even if the system is reduced.
2	The system's topology must change between a Tetra 8 and a Tetra 16.	

Case 1 is clear. In Case 2, you must make decisions:

1. Estimate your rate of future system growth.
2. Decide whether your system growth will ever require a change in topology.
3. Decide whether the present cost of buying additional hardware in a new Tetra 16 system is less than the future cost of a system outage.

Benefits of Different Topologies

The benefits of Tetra 8 topology are:

- ServerNet adapters can be put in slots that in Tetra 16 topology are reserved for SEBs and MSEBs.
- If your current needs are satisfied by a Tetra 8 topology, you do not have to purchase the extra hardware that a Tetra 16 topology requires.

The benefits of Tetra 16 topology are:

- You are not limited to the Tetra 8 maximums of eight processors and two I/O enclosures per processor enclosure.

  The number of I/O enclosures that a processor enclosure can support varies depending on the RVU of your system. Your service provider can tell you the number of I/O enclosures supported by the RVU that is running on your system.
- The system can be expanded without a system outage.

Impact on Users

For system resizing to have minimal impact on users and processes, perform online reduction and expansion tasks during periods of low use (that is, during nonpeak hours).

For Information About	See
Online reduction tasks	Section 4, Reducing a System Online
Online expansion tasks	Section 6, Expanding a System Online

System Availability

The availability of your system is determined by the communication among its:

- Enclosures
- Disks
- Processor multifunction (PMF) CRUs
- I/O multifunction (IOMF) CRUs
- IOAM enclosures (if used)
- Peripheral devices
- Subsystems

△ **Caution.** When you resize your system, you alter some or all of these paths of communication. Unless you plan your resizing carefully, you might disrupt communications in your system and reduce your system's availability.

Because the communication between and among enclosures is the basis of how your system operates, you should understand the principles of system communication:

For Information About	See
Communication paths in your system	*NonStop S-Series Planning and Configuration Guide*

System Configuration

To resize your system safely, you must know how every device in your system is configured. This knowledge will help you avoid problems such as:

- Expecting an enclosure to support more functions than it is designed for
- Removing an enclosure that supports a system-critical function

Gather all the documentation for your present system, including planning and configuration forms for all enclosures, adapters, and peripherals in your system.

If you do not have all these forms, copy forms out of the appropriate manuals and record your system configuration.

For Information About	See
Basic planning and configuration forms for your system	*NonStop S-Series Planning and Configuration Guide*
ServerNet adapters and other components installed on your system	Documentation about the ServerNet adapters and other components installed on your system
IOAM enclosures and components	*Modular I/O Installation and Configuration Guide*

Operations

When the number of enclosures in your system changes, the basic operations of the system itself are affected.

Because you are resizing your system, you must either configure added enclosures or reconfigure the system itself.

For Information About	See
Changing configurations with the Subsystem Control Facility (SCF)	*SCF Reference Manual for the Kernel Subsystem* *SCF Reference Manual for the Storage Subsystem*
Reconfiguring other subsystems and peripheral devices	Documentation that describes the ServerNet adapters and other components installed on your system

Moving a System Console

You have the choice of:

- Using one system console for the entire resizing procedure
- Using separate workstations for the donor system, the block being added, and the target system

The procedures in this guide assume the use of one system console.

Site Planning

The physical layout of your site affects how easy it is to service and operate your system.

Changing the number of enclosures in your system might involve many factors at your site including:

- Physical layout of the site
- Access to the room to install more equipment
- Cooling requirements
- Power requirements
- Space issues such as service clearance, or the capacities of cable channels and cable troughs

For Information About	See
Issues to consider when changing your site	*NonStop S-Series Planning and Configuration Guide*
Individual factors at your site	*NonStop S-Series Hardware Installation and FastPath Guide*

Cabling

When you expand a system, you must attach all cables correctly so that added enclosures can communicate with the rest of your system. When you reduce a system, you must remove cables correctly so that remaining enclosures can continue to communicate with each other.

The types of cables you need to consider are:

- ServerNet cables (discussed in this guide)
- Power-on cables (discussed in this guide)
- Ethernet cables (discussed in this guide)
- AC power cords (discussed in this guide)
- DC power cables for system enclosures with power shelves
- Cabling from SCSI ports on PMF CRUs and IOMF CRUs
- Interface cables for devices attached to ServerNet adapters
- Fibre Channel cables (to connect FCSAs to Fibre Channel storage devices such as an Enterprise Storage System (ESS))

- Emergency power-off (EPO) cables (required by some countries)

For Information About	**See**
Principles of cabling enclosures	*NonStop S-Series Planning and Configuration Guide*
How to place enclosures so that cabling is easiest	*NonStop S-Series Planning and Configuration Guide*
Attaching and removing cables	*NonStop S-Series Hardware Installation and FastPath Guide*

Other Issues

Because every system is different, this guide cannot predict every planning issue. You might want to consult your service provider.

Your consultation will be most efficient if you first collect (or fill out) all the configuration forms for your present system. With this information, you and your service provider can focus on your particular system and needs.

2 Planning System Expansion

This section defines the concept of blocks and describes how to use the worksheets in this guide to plan the orderly addition of enclosures.

About Blocks

When you expand your system online, you add enclosures in a group that is called a **block**.

Definition

A block is a group of one or more system enclosures that is supported and recognized as a unit by your system.

Whether a block is a single enclosure or multiple enclosures, it must have a properly balanced distribution of hardware and software for the topology of the system to which it will be added.

Examples

These units are blocks:

- One processor enclosure
- One I/O enclosure (or one IOAM enclosure)
- One processor enclosure with one or more I/O enclosures attached

In contrast, these units are not blocks:

- Two or more processor enclosures cabled together. This unit is not a block because each processor enclosure is a block by itself.

- Two or more I/O enclosures without a processor enclosure. This unit is not a block because if more than one I/O enclosure is in a block, the I/O enclosures must be attached to a processor enclosure.

You can think of a block as a subcomponent of your system, or part of a corner of a tetrahedron as described in the *NonStop S-Series Planning and Configuration Guide*.

Rules for Creating a Block

- Configure enclosures individually before you make them into a block.
- Cable individually configured enclosures together into blocks before adding them to your system.
- Add blocks to the system one at a time, no matter how large or small each block is.
- Add power cords and ServerNet cables to all the enclosures in a block before you add the block to your system.
- If the block includes multiple enclosures, cable all the enclosures in the block together before you add the block to your system.
- In all cases, configure an enclosure into a block offline, isolated from the system to which you will add it.

Block Structures

Your system accepts different structures of blocks, one block at a time. Therefore, you can add enclosures to your system in different groupings, depending on what you need.

For example, you might want to add one processor enclosure and one I/O enclosure to your existing system. In this situation, you have two options:

- You can add the enclosures in two different places as two blocks.
- You can combine the two enclosures into one block and then add that block to your system.

For more examples of block structures, see Examples of Block Structures on page 2-4.

The most useful way to structure the blocks for your system varies depending on your needs and the configuration of your system:

For Information About	See
The basic considerations of system topology and how it affects the number of enclosures your system contains	• NonStop S-Series Planning and Configuration Guide
How your needs dictate system topology	• Choosing a Topology on page 1-5 • Benefits of Different Topologies on page 1-6
Points to consider when choosing block structure	• Choosing Which Block Structure to Use on page 2-8

Supported Block Structures

The number of I/O enclosures that a processor enclosure can support varies depending on the RVU of your system. Thus, the number of ways a block can be structured might also vary.

For Information About	See
The number of I/O enclosures per processor supported by the RVU of your system	• *G06.nn Release Version Update Compendium* • *Interactive Upgrade Guide*
The number of I/O enclosures per processor enclosure supported by your system's topology	• NonStop S-Series Planning and Configuration Guide
How processor enclosures and I/O enclosures are configured and cabled	• NonStop S-Series Planning and Configuration Guide

Examples of Block Structures

If you want to add two processor enclosures and two I/O enclosures to your system, you can combine them into four, three, or two blocks.

Note. In the examples in this section, an I/O enclosure could instead be an IOAM enclosure. An IOAM enclosure is considered a group and can take the place of an I/O enclosure in the topology. To determine the supported group numbers for IOAM enclosures, see the *NonStop S-series Planning and Configuration Guide* or have your service provider see the *Modular I/O Installation and Configuration Guide* (in the NTL Hardware Service and Maintenance collection).

Four Blocks

You can configure each enclosure as a separate block, and add the blocks one at a time:

- Two I/O enclosures become two blocks. You add these blocks to one or two processor enclosures that already exist on the system.
- Two processor enclosures become two separate blocks.

Figure 2-1 illustrates how four enclosures can become four blocks:

Figure 2-1. Four Enclosures as Four Blocks

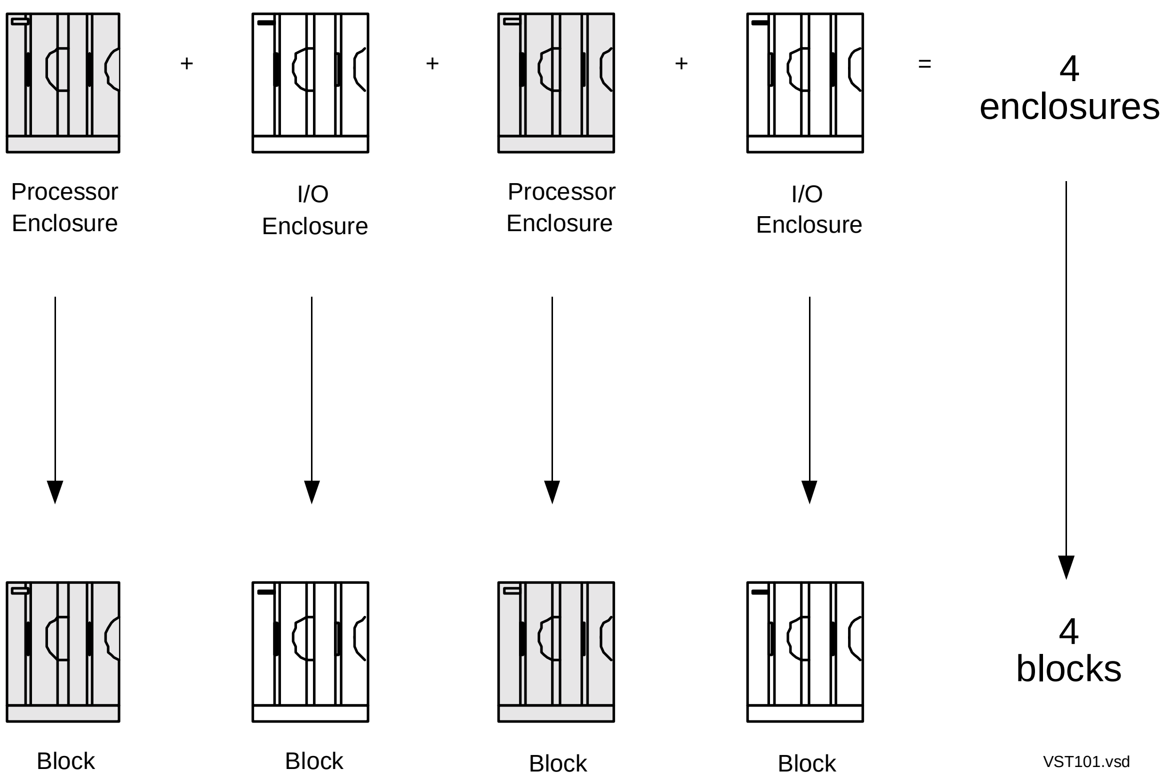

Note. Your enclosure doors might look slightly different.

Three Blocks

Your system can accept combinations of enclosures as well as single enclosures:

- One I/O enclosure becomes one block. You add this block to a processor enclosure that already exists on the system.
- One processor enclosure becomes one block.
- One processor enclosure supports one I/O enclosure. This combination becomes one block.

Figure 2-2 illustrates how four enclosures can become three blocks

Figure 2-2. Four Enclosures as Three Blocks

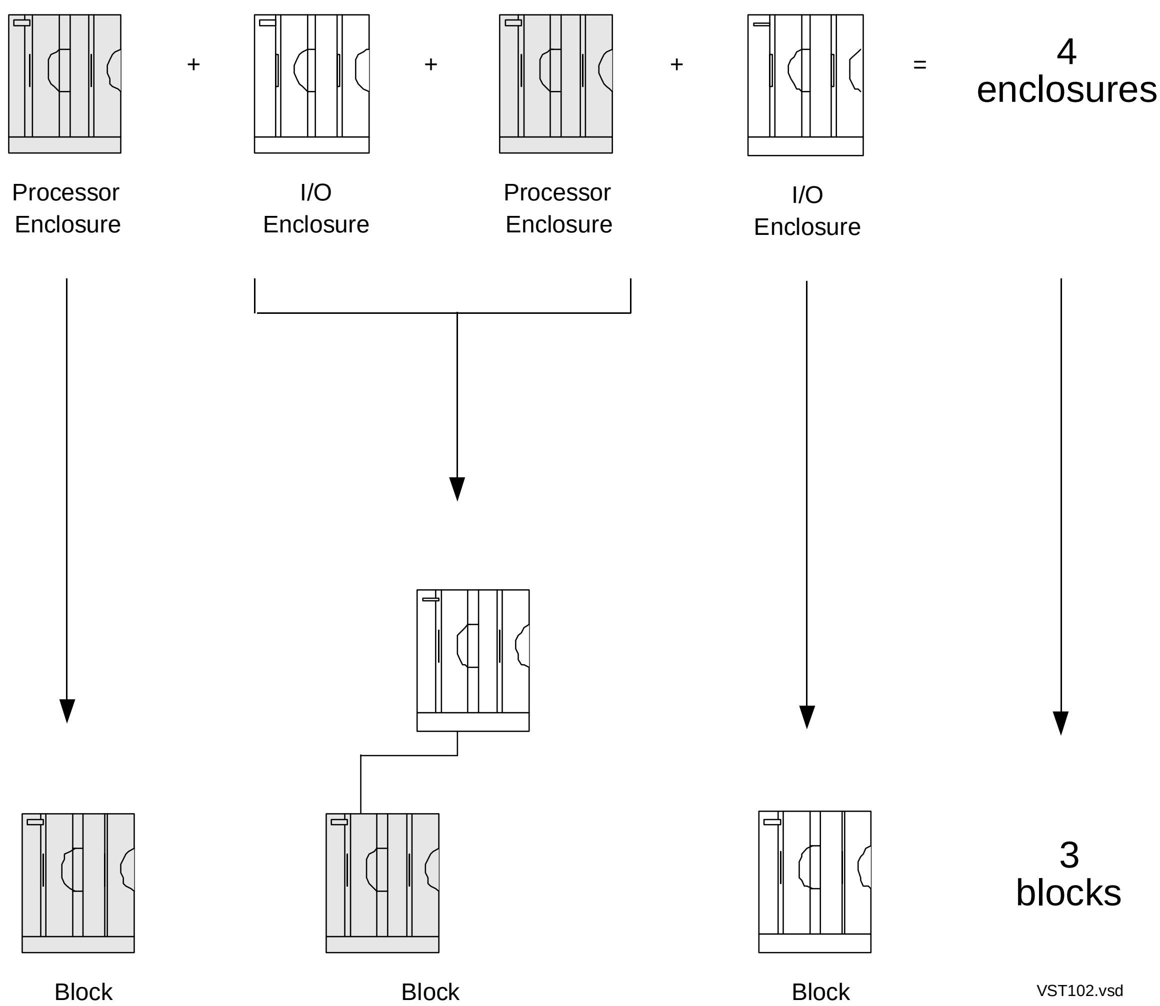

Two Blocks (Two Structures)

Your system can accept combinations of enclosures as well as single enclosures. Furthermore, these combinations can be different sizes.

Structure 1

- One processor enclosure becomes one block.
- One processor enclosure supports two I/O enclosures. This combination becomes one block.

Figure 2-3 illustrates how four enclosures can become two blocks of different sizes:

Figure 2-3. Four Enclosures as Two Blocks of Different Sizes

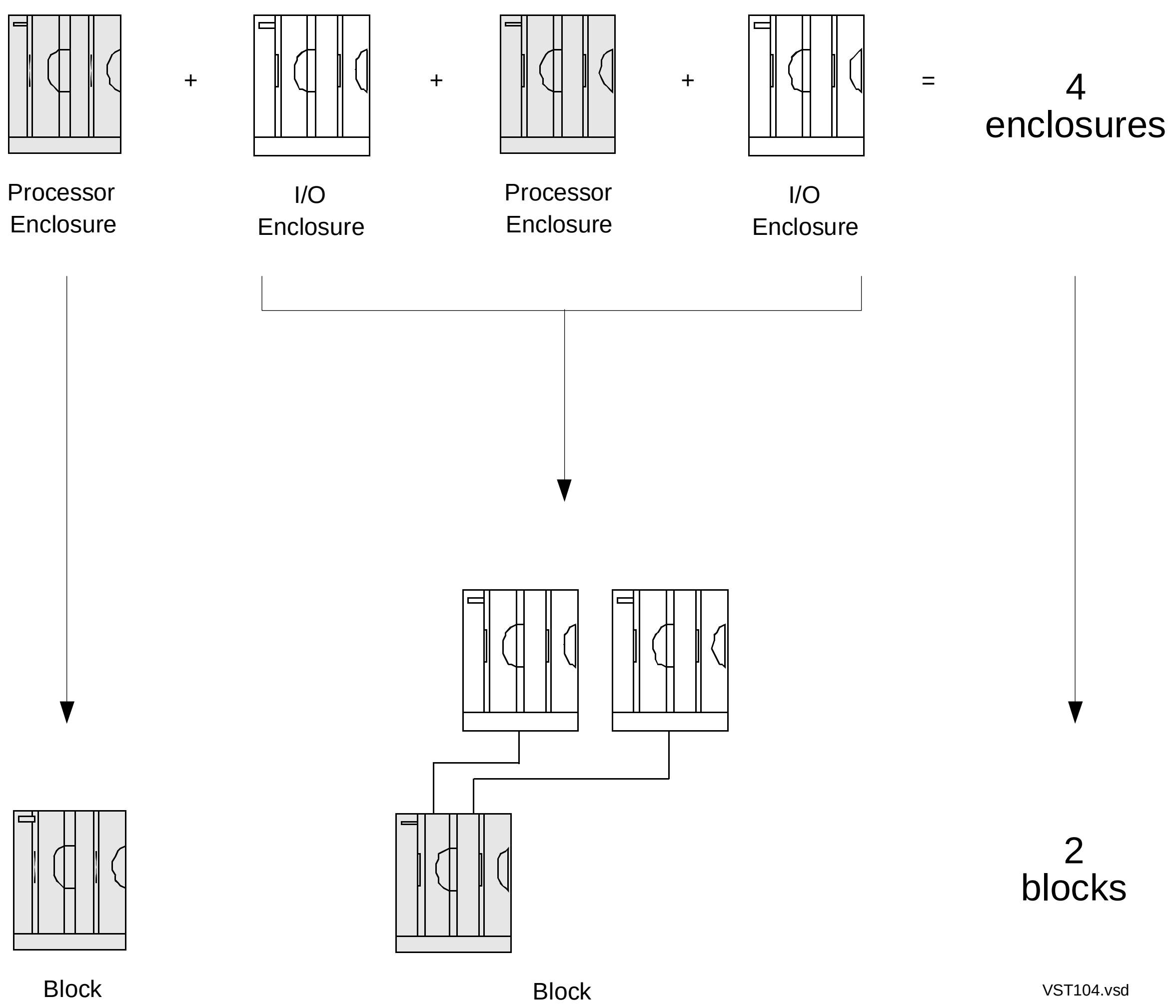

Structure 2

- One processor enclosure supports one I/O enclosure.This combination becomes one block.
- The other processor enclosure supports the other I/O enclosure. This combination becomes one block.

Figure 2-4 illustrates how four enclosures can become two identical blocks:

Figure 2-4. Four Enclosures as Two Identical Blocks

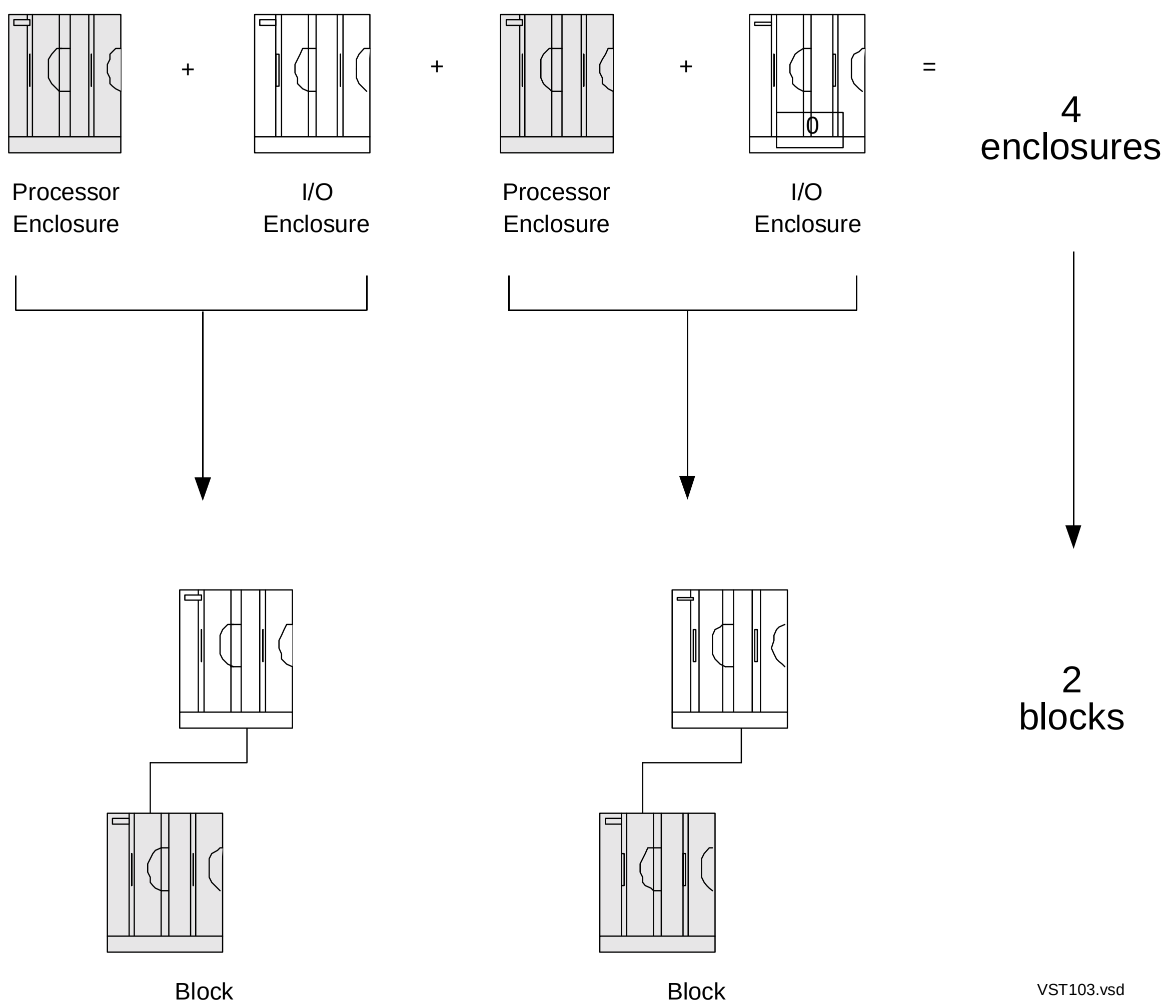

Choosing Which Block Structure to Use

Because there are so many different ways to combine enclosures into blocks, plan ahead of time which block structures will best meet your needs. Every system is different, so the best way to add blocks to your system depends on many factors:

- How you use your system
- The software and adapters on your system
- The topology of your system
- The number and arrangement within that topology of the enclosures already existing on your system
- Whether you are changing system topology as you resize your system
- Whether you might change the topology in the future
- Whether you will only expand your system or whether you might reduce it at some time as well
- The physical layout of your computer room

Issues such as whether you might change system topology in the future, or whether you will be reducing one system and using those enclosures on another system, should be planned with your systems analyst and your service provider.

However, you can clarify many of these issues by filling out Worksheets 3 and 4 (for Tetra 8 topology) or Worksheets 5 and 6 (for Tetra 16 topology) from Appendix C, Checklists and Worksheets. When you fill out these worksheets, you can see the block structures that you can use to expand or reduce your system.

The rest of this section shows how the worksheets in this guide can clarify the possibilities for resizing your system.

Example of Choosing a Block Structure

This example uses a NonStop S*xx*000 system configured as a Tetra 16 topology. The worksheet shown in this example is only for the X fabric. The tasks for the Y fabric are identical.

Example Step 1: Copy Worksheets for Your System

If you have a Tetra 16 system, locate and copy Worksheets 5 and 6 from Appendix C, Checklists and Worksheets. This figure shows an example of Worksheet 5, copied and ready to map your current system:

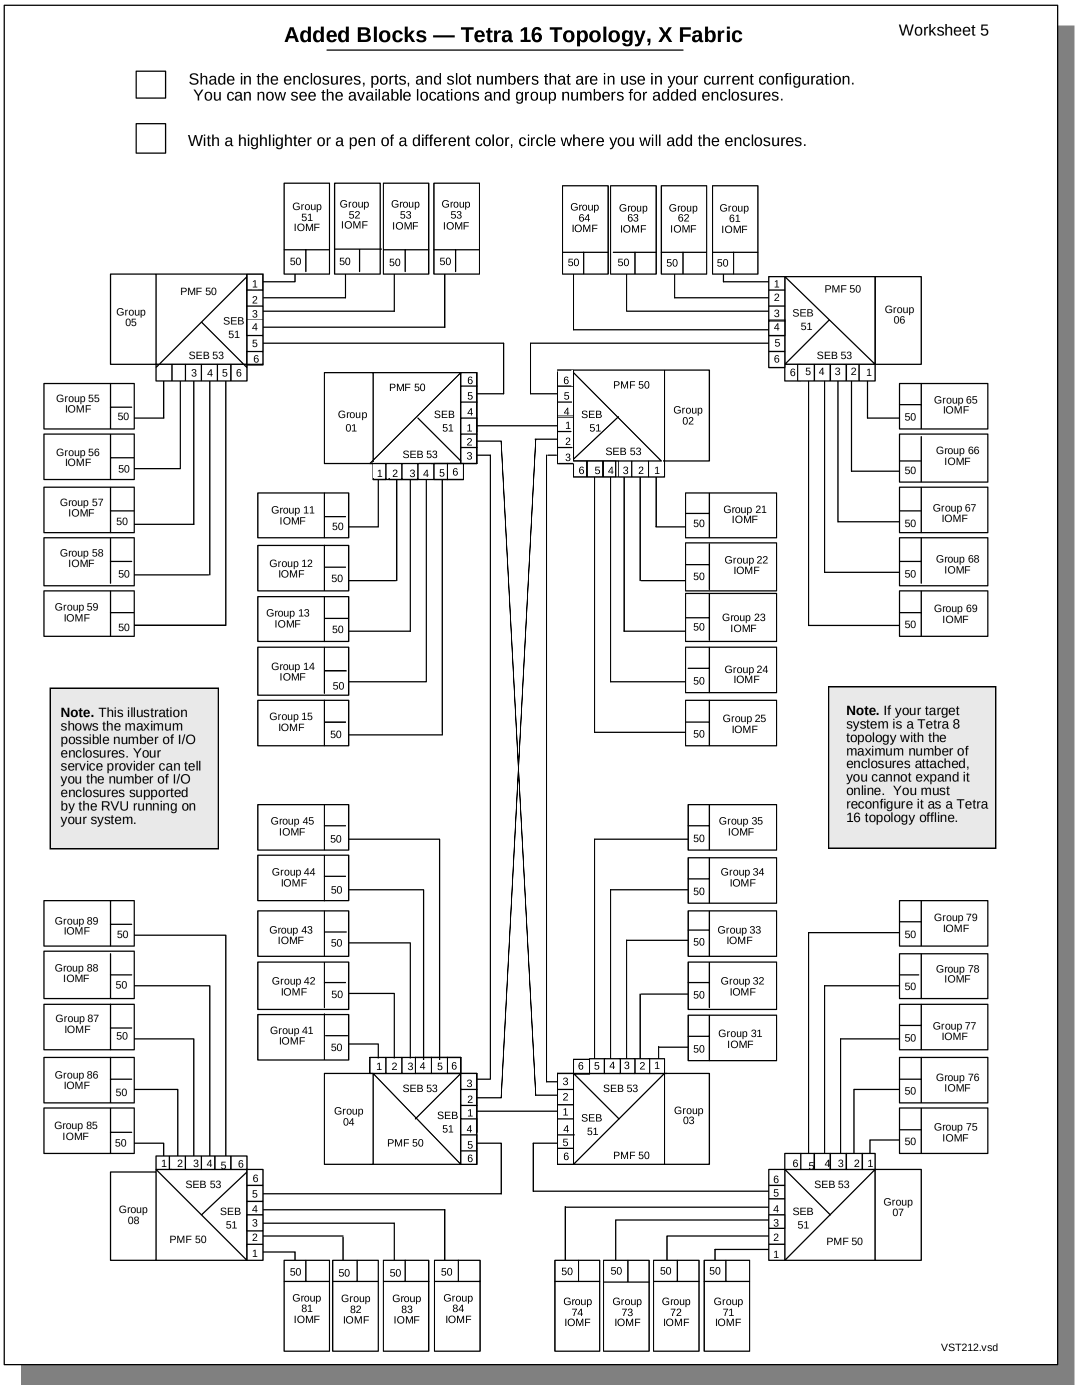

VST401.vsd

Example Step 2: Eliminate Unsupported Enclosures

Consult your service provider about how many I/O enclosures are supported for the RVU running on your system. This example assumes that you are reconfiguring a NonStop Sxx000 system that supports 36 I/O enclosures, so you mark off the unsupported I/O enclosures as shown.

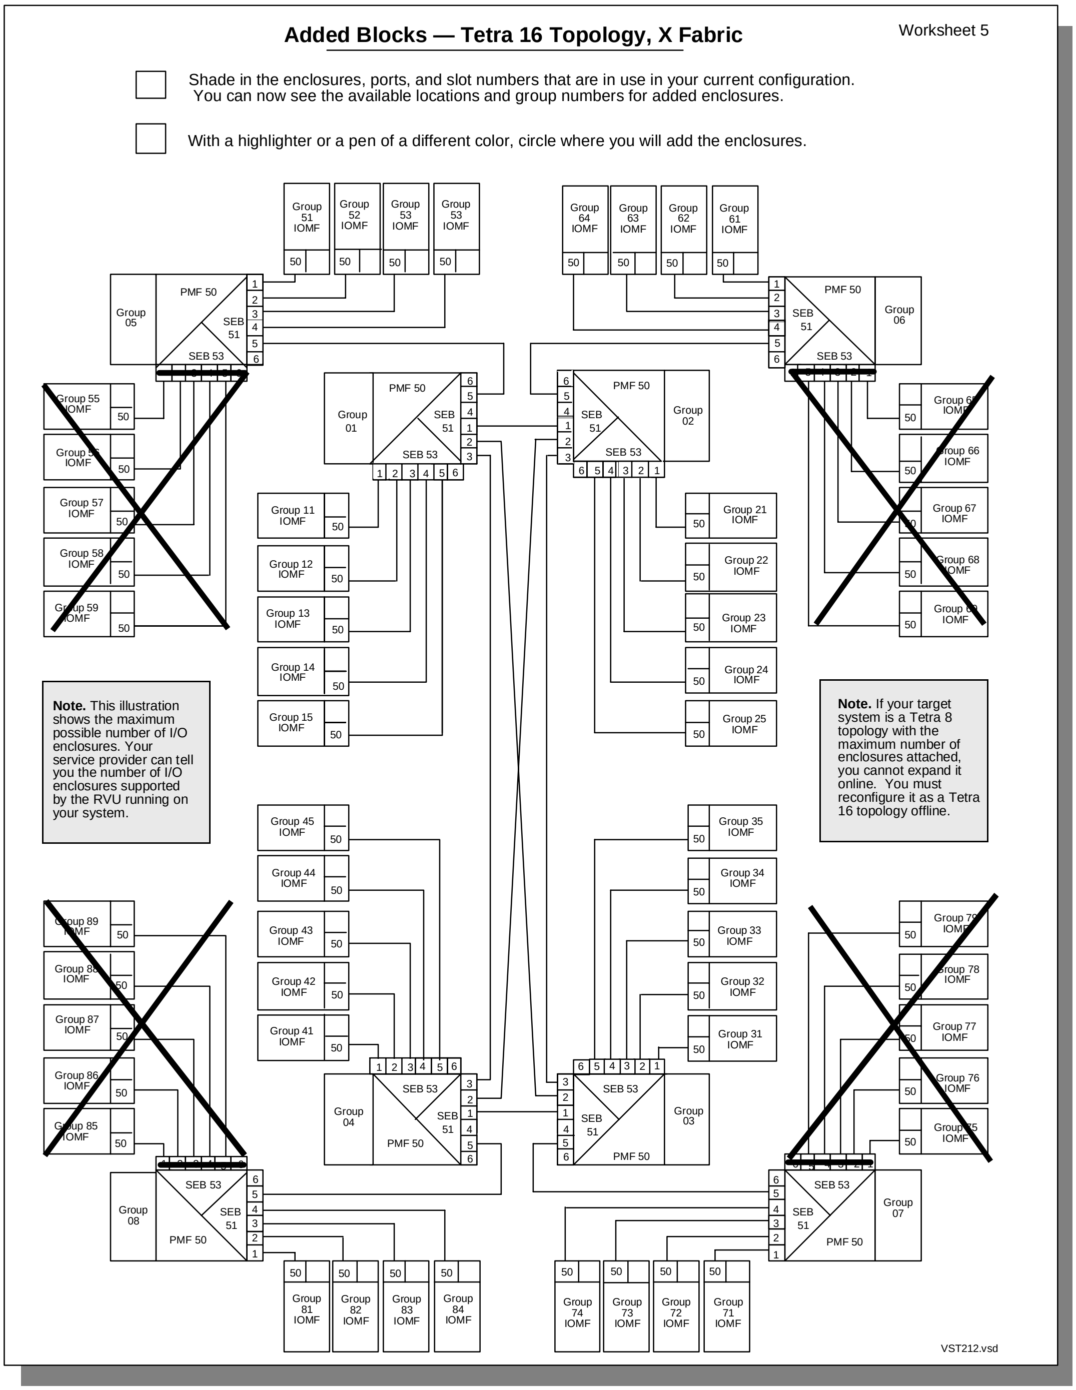

VST402.v

Example Step 3: Shade In Existing System Enclosures

Shade in the enclosures your system actually contains now. This example system contains groups 01 through 04, each with three I/O enclosures, and group 05 with two I/O enclosures. Shade in those enclosures to see the space left in your topology, and the possible block structures that fit.

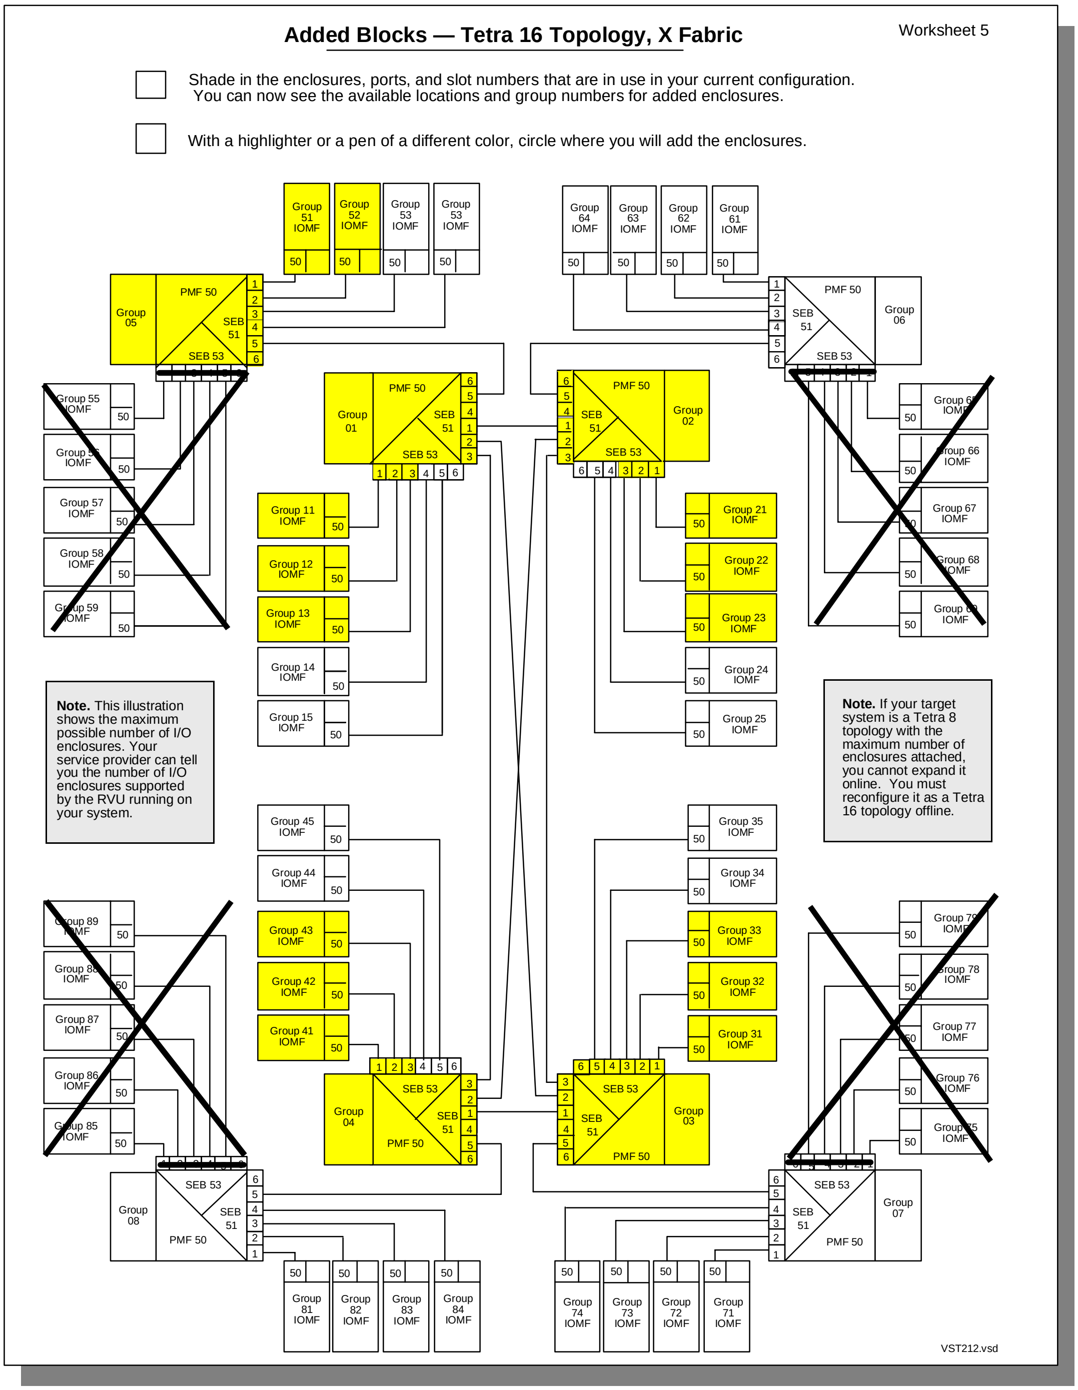

VST403.vsd

Example Step 4: Choose the Block Structure

The previous step shows that each existing processor enclosure can support two more I/O enclosures. You also can add three processor enclosures for groups 06, 07, and 08, each of which can support four I/O enclosures.

In this example, you will add two processor enclosures and five I/O enclosures. You must now decide how to build those enclosures into blocks.

The following figure shows a few of the ways you can structure these blocks. For this example, you want to balance your increased I/O capacity across as many processor enclosures as possible. So you choose a structure with five blocks, as follows:

- Three identical blocks, each made of one I/O enclosure
- Two identical blocks, each made of one processor enclosure and one I/O enclosure

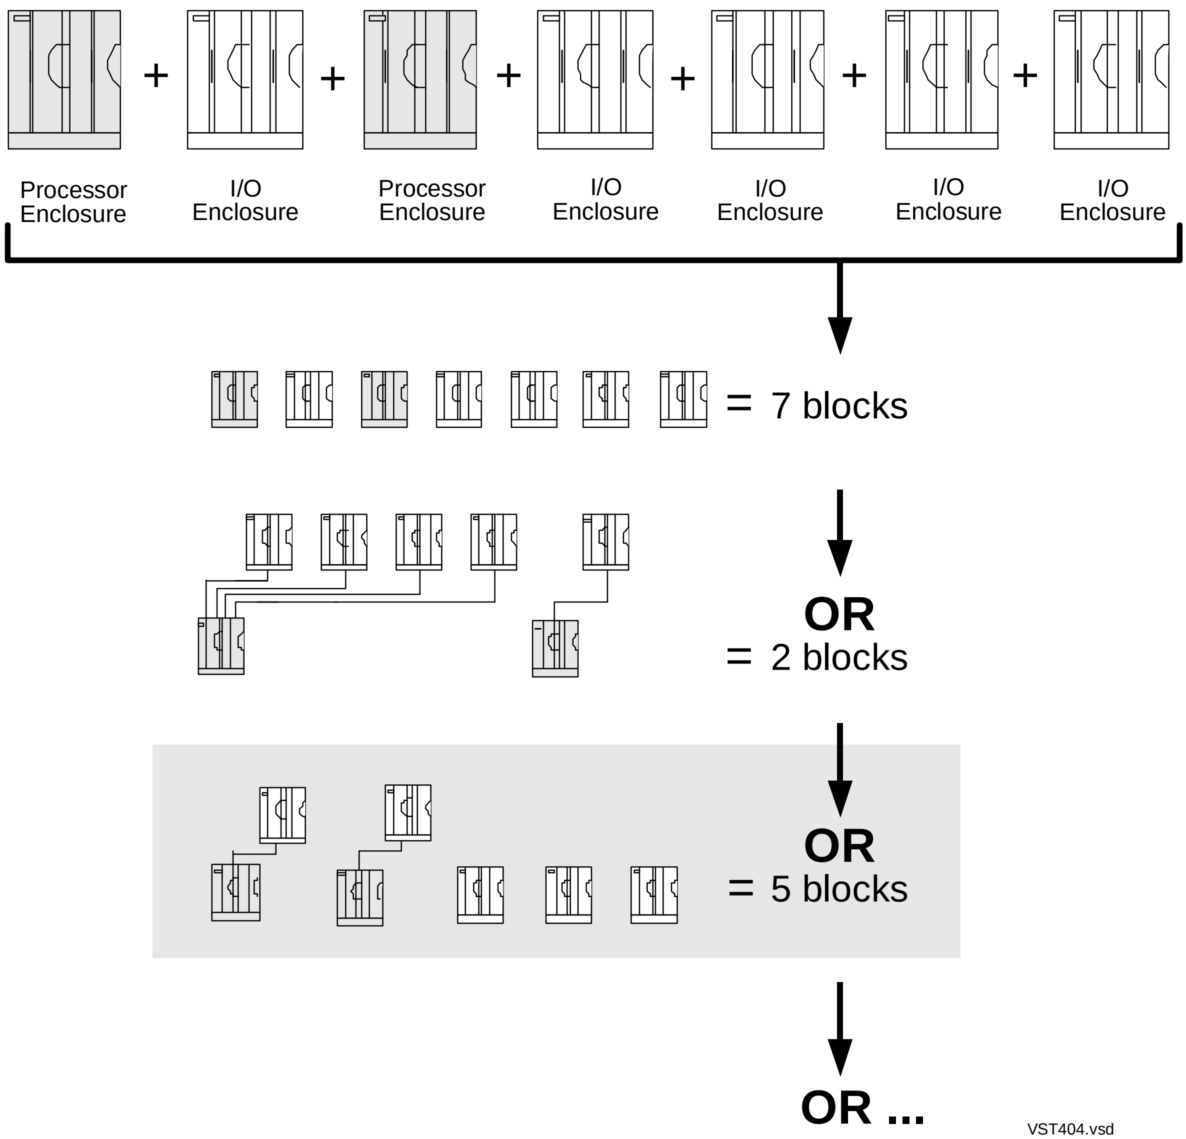

Example Step 5: Add Blocks to Worksheets

On your copy of Worksheet 5, confirm that your system will accept the block structures you have planned. Highlight the spaces in which you intend to install the blocks.

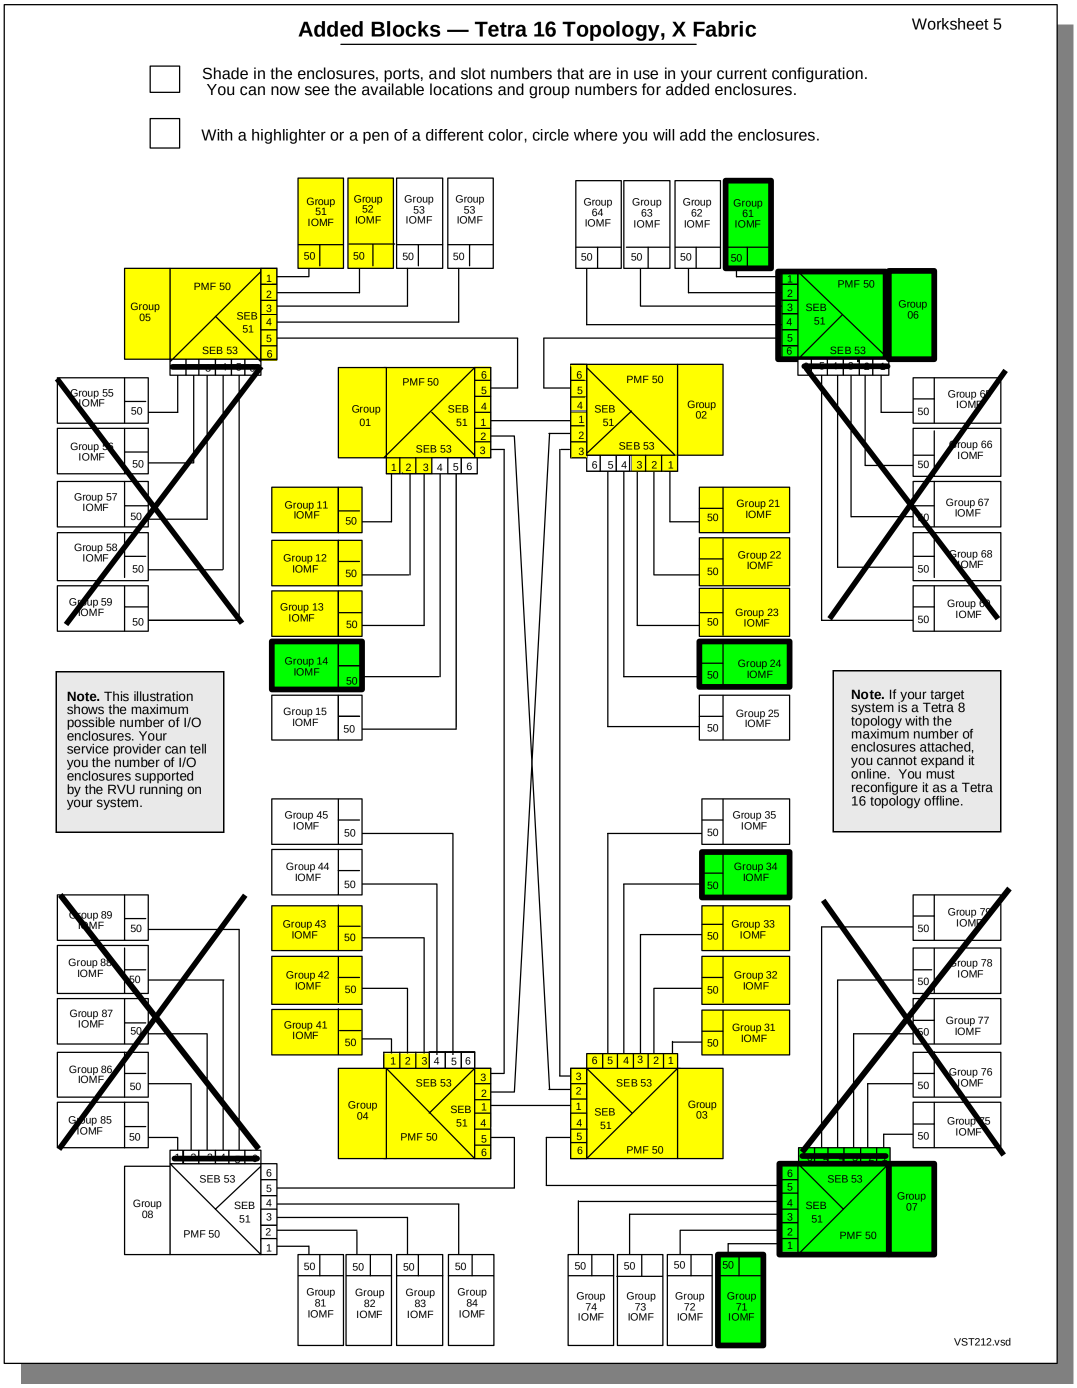

VST407.vsd

Because you have decided to add these blocks in the places shown, the blocks will be assigned the following group numbers when you attach the blocks to your system:

Block Structure	Group Numbers
One I/O enclosure	14
One I/O enclosure	24
One I/O enclosure	34
One processor enclosure attached to one I/O enclosure	06, 61
One processor enclosure attached to one I/O enclosure	07, 71

Enclosures From Another System

To expand a system online when you are using enclosures that come from another system:

Step	For more information
1. Ensure that the system you are expanding contains an adequate number of MSEBs.	Does Your System Have Enough MSEBs and MSEB PICs? on page 1-3
2. Plan the procedure of removing enclosures from another system.	Section 3, Planning System Reduction
3. Remove enclosures from the donor system.	Section 4, Reducing a System Online
4. Expand the target system.	Section 6, Expanding a System Online

To expand a system online when you are using enclosures that come directly from the factory:

Step	For more information
1. Ensure that the system you are expanding contains an adequate number of MSEBs.	Does Your System Have Enough MSEBs and MSEB PICs? on page 1-3
2. Expand your system.	Section 6, Expanding a System Online

3 Planning System Reduction

This section defines the concept of a donor system and describes how to plan the orderly removal of enclosures.

About the Donor System

Definition

You can remove enclosures from a system either to reduce that system or to add the removed enclosures to another system. The system you remove enclosures from is the **donor system**.

Caution

If any enclosures that you are planning to remove from the donor system are responsible for critical processes and resources, **you must move these processes and resources to enclosures that will remain in the system.**

Shutting down or moving processes and resources on a NonStop S-series system is a complex task with serious implications. If you are not familiar with how to shut down and move devices or subsystems in your system, see the appropriate manuals.

Considerations

Before you remove any enclosures, answer the following questions about the donor system:

Are Both ServerNet Fabrics Up?

For online reduction to succeed, both ServerNet fabrics in the system must be fully operational. If any ServerNet alarms have been generated, you must correct these problems before you reduce the system online.

Can an Enclosure Be Removed Easily?

△ **Caution.** Enclosures are not customer-replaceable units (CRUs). Removing enclosures should be done only by service providers trained by HP.

To be easily removed, an enclosure should be either in a single-high stack or on the top of a double-high stack.

Removing an enclosure from the bottom of a double-high stack requires extra people and clearance space to:

- Lift the top enclosure far enough that the bottom enclosure can be rolled out from beneath it
- Ensure that, as the enclosures are rolled apart, no cables or cords are twisted, crushed, or detached
- Lower the top enclosure safely to the floor, onto another base enclosure, or onto a frame base

The number of extra people and clearance space you need depends on the hardware components in the donor system and on the site layout:

For Information About	See
Procedures for removing enclosures	*NonStop S-Series Service Provider Supplement* (available to service providers)

Remove All Enclosures at Once or One at a Time?

The two methods of removing enclosures each have advantages and disadvantages.

If you are removing only one enclosure, there is no difference in these methods.

Removing All Enclosures At Once

Procedure

1. Route all applications and processes away from all enclosures to be removed.
2. Delete all the enclosures from the system configuration at once.
3. Remove them physically from the system one at a time, as described in Section 4, Reducing a System Online.

Advantage

This method allows you to reconfigure your donor system only once, rather than after each enclosure has been prepared and removed. Therefore, your system preparation time is significantly reduced.

Disadvantage

This method can increase the difficulty of pinpointing the source of any problems.

Removing Enclosures One at a Time

Procedure

1. Route all applications and processes away from the enclosure to be removed.
2. Delete the enclosure from the system configuration.
3. Remove the enclosure physically from the system.
4. Repeat this procedure for each enclosure to be removed.

Advantage

This method makes it easier to pinpoint the source of any problems.

Disadvantage

System preparation time is increased because you reroute all applications and processes and reconfigure the donor system one enclosure at a time.

Is the Enclosure Providing Critical Resources?

Information about critical resources is system-specific. Therefore, consult your system analyst and service provider to determine the actions most appropriate for your system.

Determine the following information about each enclosure you will remove:

- What devices, adapters, and processes it contains
- What communication paths it maintains for other enclosures

You need to know this information so that you can reconfigure your system to maintain any processes and communication paths that otherwise would be disrupted when enclosures are removed.

For example:

- Removing an enclosure that contains Ethernet 4 ServerNet Adapters (E4SAs), Gigabit Ethernet ServerNet Adapters (GESAs), or Gigabit Ethernet 4-Port ServerNet Adapters (G4SAs) will disrupt TCP/IP processes that are configured to use the Ethernet adapters.
- Removing an enclosure that has Kernel-Managed Swap Facility (KMSF) swap files on its disks might cause applications or system processes to terminate.

For Information About	See
KMSF swap files	*Kernel-Managed Swap Facility (KMSF) Manual*

Devices, Adapters, and Processes

Devices, adapters, and processes in the enclosure you plan to remove can be located by using either the OSM Service Connection or TSM Service Application.

Before you begin the reduction procedure, record this information on the appropriate worksheets in Appendix C. Then, with your system analyst, determine critical operations that should first be moved to other enclosures or terminated.

Indirect Communication Paths

Neither the OSM Service Connection or TSM Service Application displays indirect paths of communication. However, these indirect paths can affect other enclosures significantly.

For example: groups 08 and 41 communicate with each other through group 04. Group 41 might connect through group 04 to use a process that is running in group 08. Therefore, if you remove group 08, group 41 is also affected.

Dependent processes and indirect communication paths can be of three types:

- Active processes and communication paths using a device
- Active processes and communication paths not currently using a device

- Processes and communication paths not started that might need a device

For more information, see Determining Communication Paths on page 3-6.

Is the Enclosure an IOAM Enclosure?

If you are planning to remove an OAM enclosure from your system, the enclosure must be removed by a service provider trained by HP. See Expansion With IOAM Enclosures on page 1-2.

Determining Communication Paths

You must determine the direct and indirect communication paths within your donor system so that you know which enclosures to prepare for online reduction. To trace the ServerNet connections among enclosures, see Worksheets 9, 10, and 11 in Appendix C.

When you see from these worksheets which enclosures might be affected by system reduction, determine with your system analyst or service provider what processes and other devices might also be affected.

Because the adapters and communications devices in each system are configured uniquely to that system, these worksheets cannot show every possible direct or indirect connection. For example, SWAN concentrators might be connected to your system in several places.

For Information About	See
The definition of an indirect communication path	Indirect Communication Paths on page 3-4
ServerNet connection paths in detail	*NonStop S-Series Planning and Configuration Guide*

Note. The cabling diagrams in this guide identify all ServerNet expansion boards as SEBs. Your system might have or require modular ServerNet expansion boards (MSEBs) in the slots designated for SEBs. SEBs and MSEBs are functionally equivalent.

Example Using the ServerNet Worksheets

This subsection illustrates the following steps in completing a ServerNet worksheet:

These worksheets illustrate the maximum configuration of each topology. Your system might contain fewer enclosures. You must note what your system currently contains to trace the pathways of ServerNet communication accurately.

1. Choose and Copy the Appropriate Worksheets

If the system topology is …	Copy these worksheets from Appendix C
Tetra 8	ServerNet Worksheet 9
Tetra 16	ServerNet Worksheets 10 and 11

This figure shows Worksheets 10 and 11 for a Tetra 16 system.

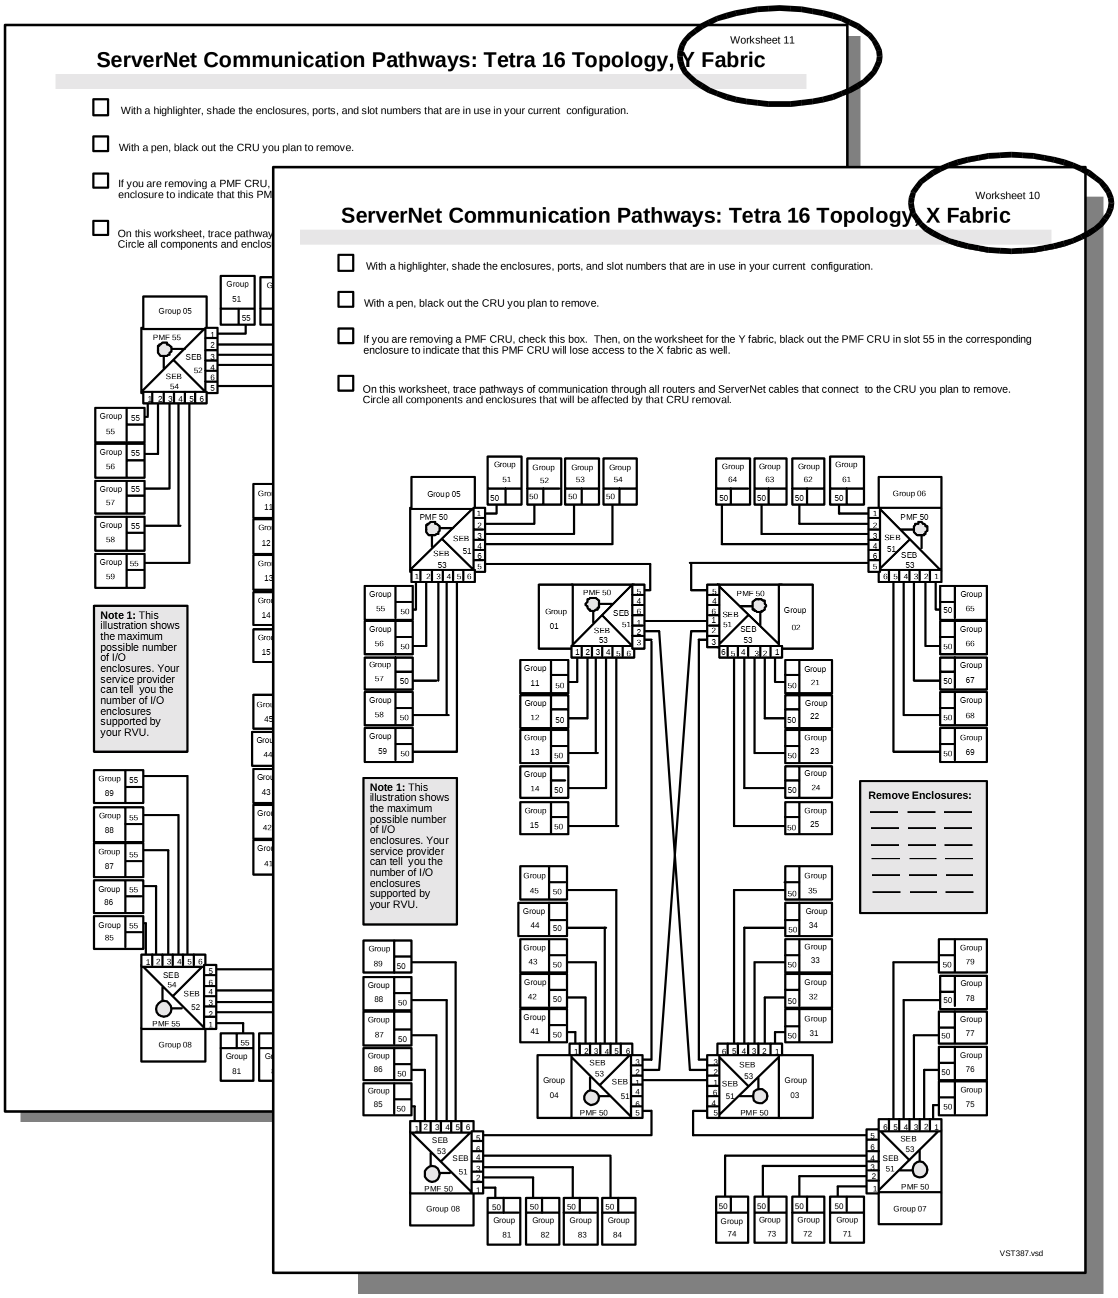

VST392.vsd

2. Highlight Your Present System Configuration

Using a highlighter on a copy of Worksheet 9 or Worksheets 10 and 11, shade in all the ports, slots, and enclosures your current system contains.

The next example worksheet shows a system that contains the following enclosures:

Processor Enclosure Groups	I/O Enclosure Groups
01	11, 12, 13, 14, 15
02	21, 22, 23, 24, 25
03	31, 32, 33, 34, 35
04	41, 42, 43, 44, 45
05	51, 52, 53, 54, 55, 56, 57
06	61, 62, 63, 64, 65, 66

△ **Caution.** The connections to IOAM enclosures are different from those shown in the following topology diagrams. A service provider trained by HP must remove IOAM enclosures from your system.

Configured portions of the system are shaded, while unconfigured portions of the system are white:

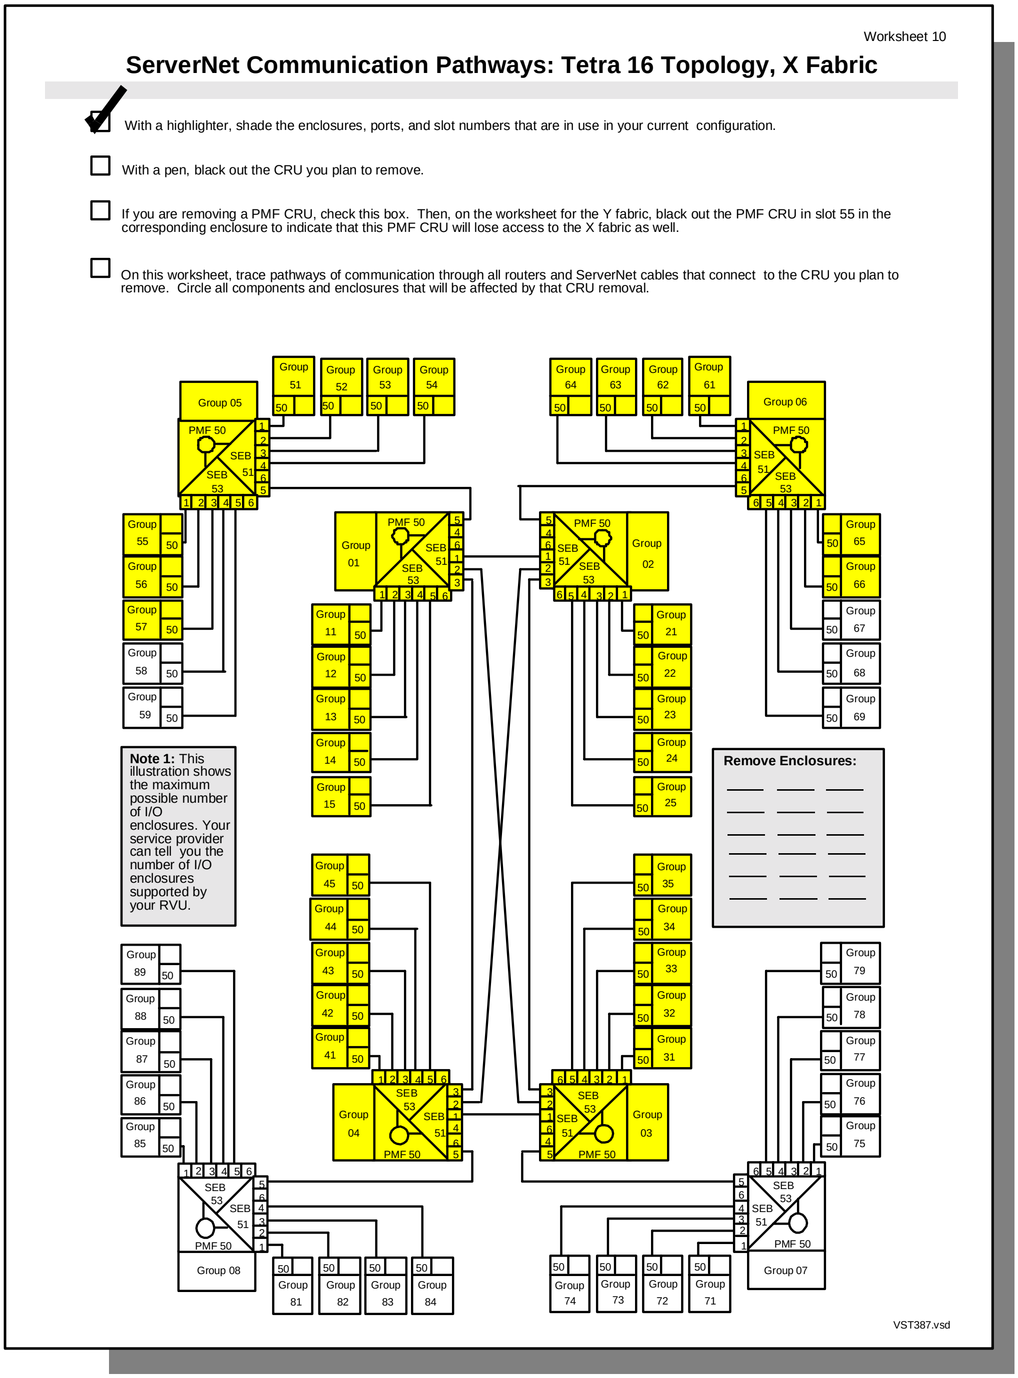

VST393.vsd

3. Black Out the Enclosures to Be Removed

With a marker, darken the enclosures you want to remove.

This example worksheet shows a Tetra 16 system in which group 02 will be removed:

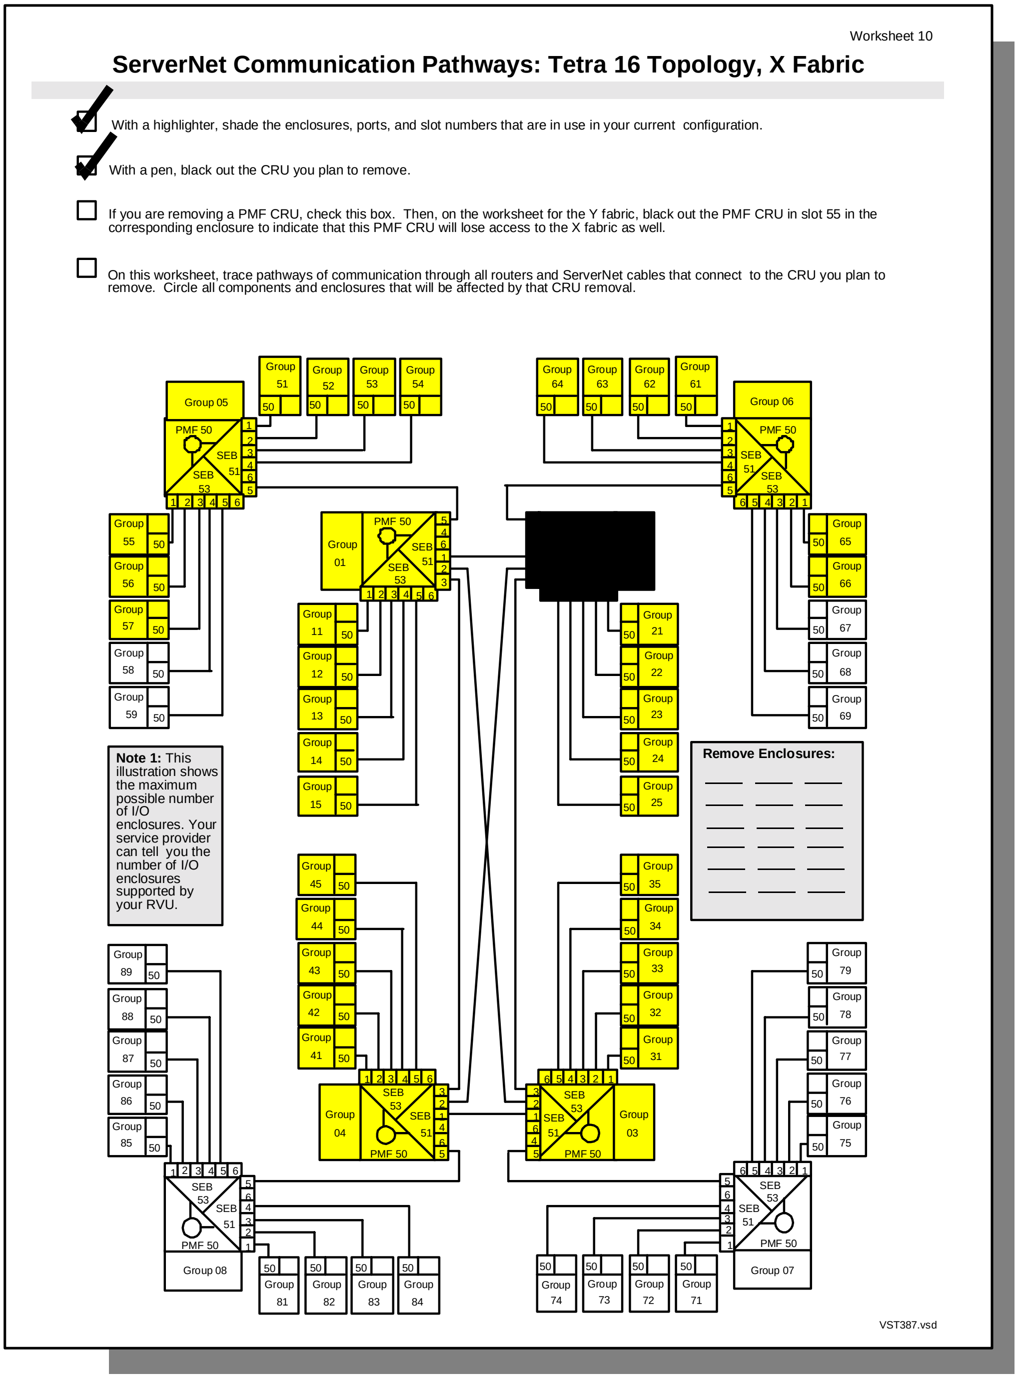

VST394.vsd

4. Trace Communication Paths to Each Enclosure

Removing an enclosure removes communication paths within your system. In this step, you determine what communication paths will be removed from your system:

1. Starting from the central “X” in the tetrahedral topology for the system, trace all ServerNet cables to the ServerNet expansion boards (SEBs) or modular SEBs (MSEBs) in the processor enclosures.
2. Trace through those SEBs or MSEBs to the attached enclosures.
3. Do not trace through the blacked-out enclosure. All communication paths end when they reach that enclosure.

The following example worksheet shows a Tetra 16 system from which you want to remove group 02. Because group 02 is blacked out, you cannot trace any paths through the SEBs or MSEBs in group 02 to reach any other enclosures attached to group 02. Therefore, no paths reach group 06.

Worksheet 10

ServerNet Communication Pathways: Tetra 16 Topology, X Fabric

- [x] With a highlighter, shade the enclosures, ports, and slot numbers that are in use in your current configuration.
- [x] With a pen, black out the CRU you plan to remove.
- [x] If you are removing a PMF CRU, check this box. Then, on the worksheet for the Y fabric, black out the PMF CRU in slot 55 in the corresponding enclosure to indicate that this PMF CRU will lose access to the X fabric as well.
- [x] On this worksheet, trace pathways of communication through all routers and ServerNet cables that connect to the CRU you plan to remove. Circle all components and enclosures that will be affected by that CRU removal.

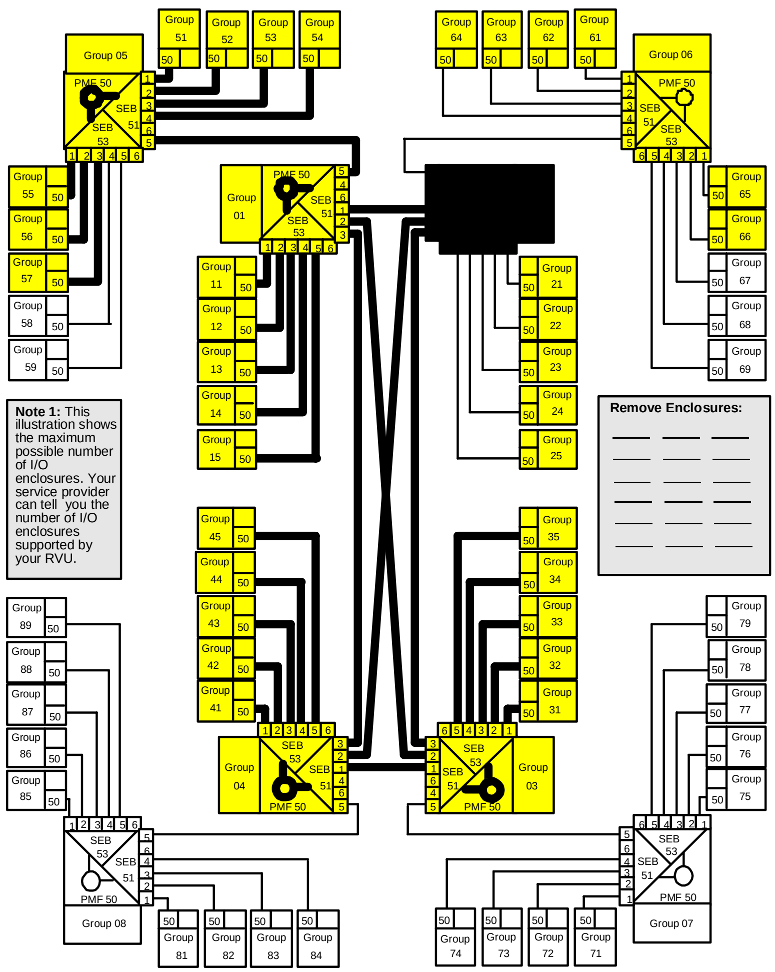

VST387.vsd

VST395.vsd

5. Circle the Affected Enclosures

After you have traced all possible communication paths, you can see the enclosures that will be affected by enclosure removal. Circle those enclosures. You must prepare to remove these parts of your system.

△ **Caution.** If your tracing of communication paths is incomplete or incorrect, the procedure you follow in Section 4, Reducing a System Online, might remove more enclosures than you intended.

The following example worksheet shows that you circle group 02 because you are planning to remove it. Because group 02 has I/O enclosures connected to it, circle those I/O enclosures as well. Group 02 also provides the communication paths to group 06 and its associated I/O enclosures, so circle those enclosures as well.

The enclosures you must prepare for the removal procedure are:

Processor Enclosure Groups	I/O Enclosure Groups
02	21, 22, 23, 24, 25
06	61, 62, 63, 64, 65, 66

Stop any system-critical processes or devices that are located in those enclosures, or move these processes and devices to enclosures that will remain in the system.

Worksheet 10

ServerNet Communication Pathways: Tetra 16 Topology, X Fabric

- [x] With a highlighter, shade the enclosures, ports, and slot numbers that are in use in your current configuration.
- [x] With a pen, black out the CRU you plan to remove.
- [x] If you are removing a PMF CRU, check this box. Then, on the worksheet for the Y fabric, black out the PMF CRU in slot 55 in the corresponding enclosure to indicate that this PMF CRU will lose access to the X fabric as well.
- [x] On this worksheet, trace pathways of communication through all routers and ServerNet cables that connect to the CRU you plan to remove. Circle all components and enclosures that will be affected by that CRU removal.

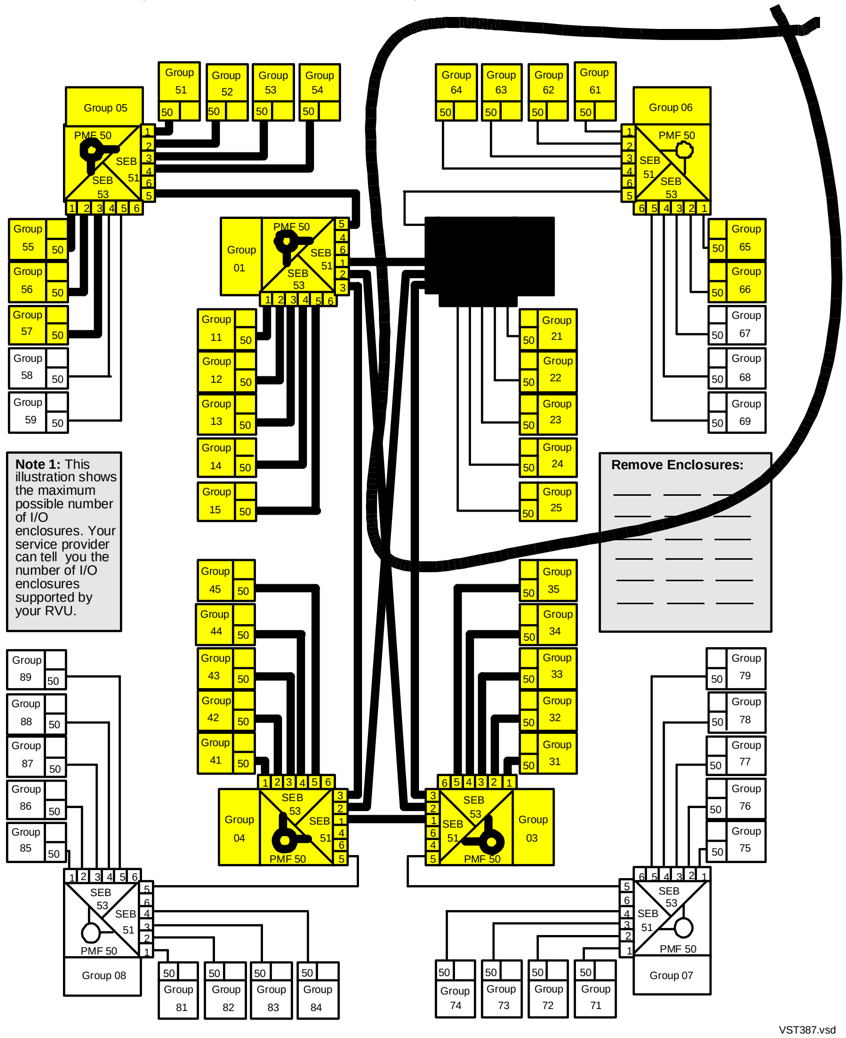

VST396.vsd

6. List Affected Enclosures in Reverse Numerical Order

The following example worksheet shows that you list the enclosures you must prepare for removal. These group numbers are listed in the Enclosures to be Removed box:

66, 65, 64, 63, 62, 61, 25, 24, 23, 22, 21, 06, 02

Physically remove the enclosures from your system in this order. For more information, see Order of Removal on page 3-17.

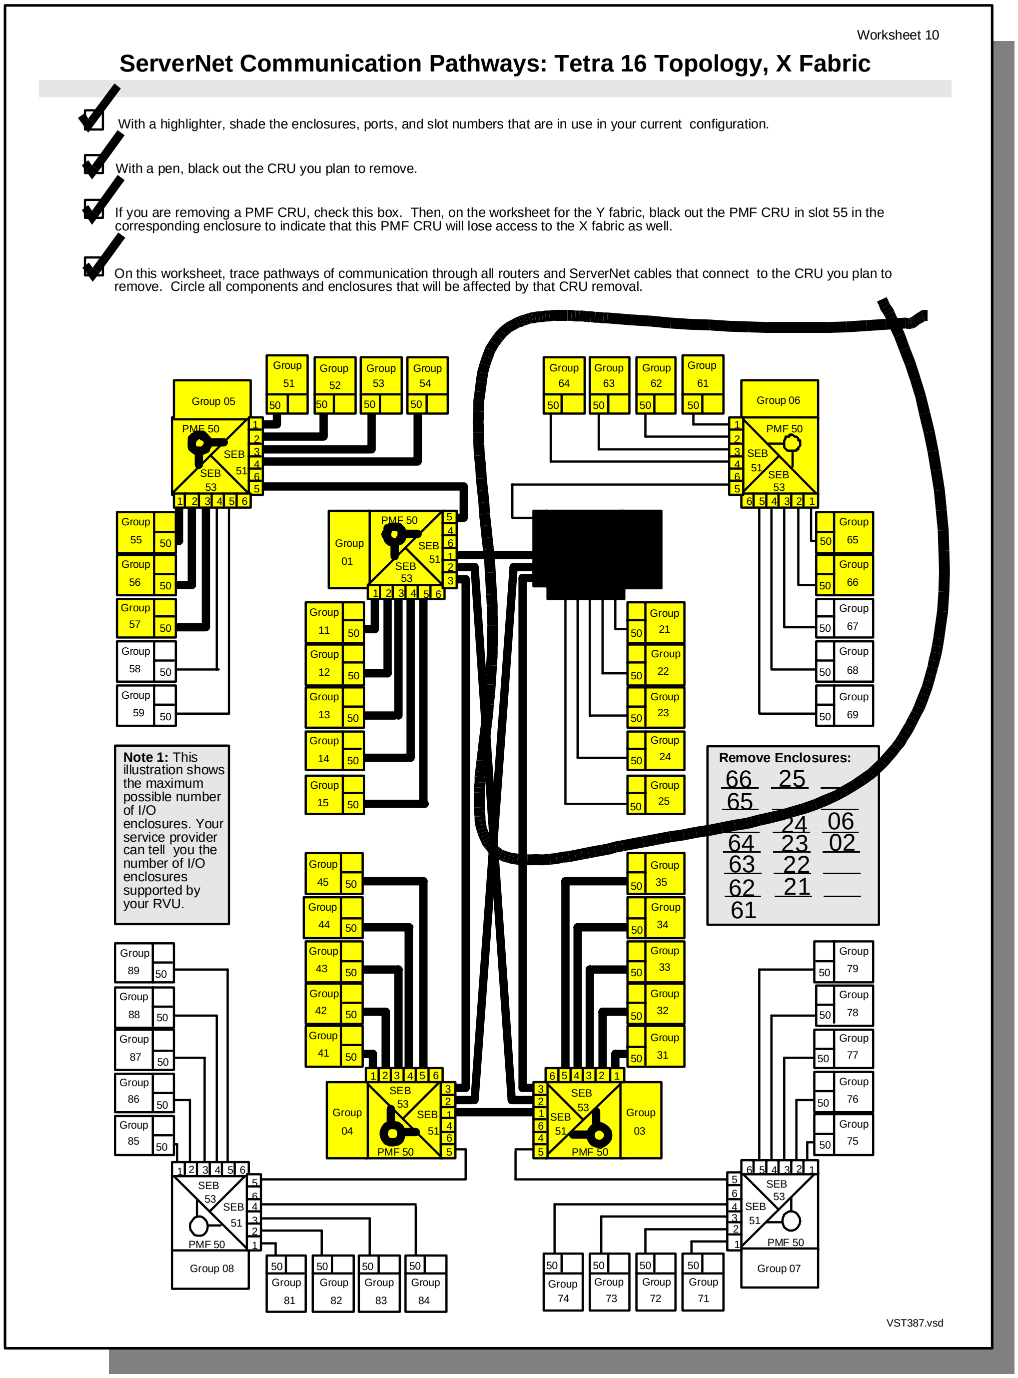

VST397.vsd

Order of Removal

After you have prepared your system, power off and remove enclosures starting with the largest enclosure number. Proceed in decreasing numerical order by enclosure number, and end with the smallest enclosure number.

For example, if you remove group 02, you affect the associated I/O enclosure groups 21, 22, 23, 24, and 25. You also affect group 06 and its associated I/O enclosure groups 61, 62, 63, 64, 65, and 66.

Therefore, remove these enclosures from your system one at a time in this order:

66, 65, 64, 63, 62, 61, 25, 24, 23, 22, 21, 06, 02

△ **Caution.** Never remove group 01. Doing so will halt your system.

The Next Step

When you finish planning which enclosures to remove from the donor system, proceed to Section 4, Reducing a System Online.

4 Reducing a System Online

This section describes the procedure for removing enclosures from a system online, either to reduce a system or to expand a target system.

Note. You must perform each step as documented, in the order shown.

You should also read and understand the discussions and concepts in Section 3, Planning System Reduction, before reducing a system online.

For information about offline reduction, see Section 5, Reducing a System Offline.

△ **Cautions.**

- You can reduce your system online only if it is running G06.04 or later, with a compatible TSM package or G06.08 or later, with a compatible OSM or TSM package. If you are running an earlier RVU than G06.04, you must reduce the system offline. See Section 5, Reducing a System Offline.
- Never remove group 01. Doing so will halt your system.
- If any enclosures that you plan to remove are responsible for critical processes and resources, first reconfigure your system to redirect these processes and resources to enclosures that will remain in the system or terminate the processes and resources.
- Make sure you have KMSF swap files available for the resources remaining on the donor system after the reduction is finished. Otherwise, after you resize your system, applications and processes on the resized system might terminate. For more information, see the *Kernel-Managed Swap Facility (KMSF) Manual*.
- Before you physically remove an enclosure, record all information such as IP addresses and firmware versions, as described in this section. After an enclosure is disconnected from the system, this information is inaccessible.
- Even if you are sure you already know this information, complete the appropriate worksheets in Appendix C. If you do not record this information correctly now, you might not be able to configure these enclosures when they are added elsewhere later.
- Before removing enclosures, complete Worksheets 1 and 2.*n* for each enclosure.

1. Prepare the Donor System for Reduction

1.1 Prepare Worksheets

1. From Appendix C, make copies of Worksheet 1 and Worksheets 2.1 through 2.8 (referred to hereafter as Worksheets 2.*n*).
   If the system contains any IOAM enclosures, also make copies of Worksheet 2.1a.
2. Circle the words *Donor System* at the top of each worksheet.
3. Make one copy of Worksheet 8 in Appendix C for each enclosure you will remove. If any of the enclosures is an IOAM enclosure, make a copy of Worksheet 8a.
4. Staple these worksheets together into a **donor system** packet.

If you will be adding these enclosures to another (target) system, be sure to keep this donor packet separate from worksheets for the target system. It is important not to mix information from your target system with information from your donor system.

1.2 Upgrade OSM or TSM Software If Necessary

OSM Client Software

To determine whether the OSM Client Software components need to be upgraded:

Step	Action	Use This Procedure
1.	Determine the current version of the OSM Client software components.	Determine the Product Versions of the OSM Client Software on page A-7.
2.	Determine whether you are running the latest versions of the client software.	See the *OSM Migration Guide*
3.	Upgrade the client software, if necessary	See the *NonStop System Console Installer Guide*.

OSM Console Tools

HP recommends that you install the OSM Console Tools to get the most recent versions of the OSM Service Connection and the OSM Event Viewer. See the *NonStop System Console Installer Guide*.

TSM Software

If you are reducing your system online, you must be running both the G06.04 or later RVU and a compatible version of the TSM package. If you are running an earlier RVU than G06.04, you must reduce the system offline, as described in Section 5, Reducing a System Offline.

Step	Action	Use This Procedure	Record Information Here
1.	Record the current OSM or TSM version.	Determine the Product Version of the TSM Client Software on page A-7.	**Worksheet:** 1 **Section:** TSM client software on donor system console, Initial Version.
2.	Verify that you are running a version of the TSM client software that is compatible with the TSM server software on the system you plan to connect to.	**Start > Programs > Compaq TSM > TSM Documentation > TSM Read Me > TSM Product Version Information > TSM Software Product Versions**	**Worksheet:** 1 **Section:** TSM client software on donor system console, Needs to Be Upgraded to
3.	If the two TSM versions differ, upgrade the TSM client software.	*NonStop System Console Installer Guide*	
4.	Record the upgrade.		**Worksheet:** 1 **Section:** TSM client software on donor system console **Upgraded** checkbox

1.3 Save the Current Donor System Configuration

Before you alter the donor system's configuration, save its current configuration for fallback purposes:

1. Log on to TACL, start SCF, and enter the SAVE CONFIGURATION command:

   ```
   >SCF
   -> SAVE CONFIGURATION xx.yy
   ```

   where *xx.yy* is a configuration number in the range 00.00 through 99.99 that is not already in use.

   If you specify a number that is already in use, you receive this message:

   ```
   Configuration file $SYSTEM.ZSYSCONF.CONFxx.yy already exists.
   Replace it (Y/[N])
   ```

2. Enter `N` and retry the SAVE CONFIGURATION *xx.yy* command using another number.

3. Record the number of your saved configuration (*xx.yy*) on Worksheet 1.

2. Record Information About the Donor System

Before you physically remove enclosures from your donor system, you must record information about those enclosures on the worksheets in Appendix C. You determine this information by using a system console.

2.1 Record Basic Information

Continue filling out worksheets for the donor system, as follows:

Step	Action	Result
1.	On the OSM or TSM Low-Level Link tool bar, click **System Discovery**.	After several minutes, the Management window appears. The Attributes tab displays the following: ● System Name ● System Serial Number ● System Type ● Topology ● SP Firmware Version (on status bar at bottom of main window)
2.	Record this information on Worksheet 1 of the donor system packet.	

If you are performing offline system reduction, return to Section 5, Reducing a System Offline.

2.2 Determine IP Addresses and Firmware Versions

Perform this step twice for each enclosure you are removing from the donor system:

- Once for the PMF CRU or IOMF CRU in slot 50
- Once for the PMF CRU or IOMF CRU in slot 55
- For an IOAM enclosure, once for each ServerNet Switch Board

Step	Action	Result
1.	In the tree pane, located on the left side of the window, expand the icon for the system.	The tree pane displays all the enclosures.
2.	Expand the group and module so that you can see the PMF or IOMF CRUs. For an IOAM enclosure, expand the group, the IOAM enclosure, and the modules so that you can see the ServerNet Switch Boards	The Physical view displays the devices in the enclosure.
3.	Expand the PMF or IOMF in slot 50 and click on the SP. For an IOAM Enclosure, click on the ServerNet Switch Board in module 2.	The attributes appear in the detail pane, including: • IP Address • SP Firmware Version If this is a processor enclosure, you also see: • Processor Firmware Version • SCSI Firmware Version • The Topology is either 8 (for Tetra 8) or 16 (for Tetra 16). If this is a ServerNet Switch board, you see: • ME Firmware version • ME FPGA version
4.	Record this information on Worksheet 2.*n* of the donor system packet. For IOAM enclosures, use Worksheet 2.1a of the donor system packet.	
5.	For the CRU in slot 55, repeat steps 3 and 4. For a ServerNet Switch Board in module 3 of an IOAM enclosure, steps 3 and 4.	
6.	For every enclosure you are removing from the donor system, repeat steps 2 through 5. For information about which enclosure to remove next, see Order of Removal on page 3-17.	

Note. IP addresses for the ServerNet Switch boards in an IOAM enclosure are dynamically assigned by a DHCP server or software that performs the same function.

2.3 Label Enclosures

You must identify each enclosure you plan to remove before you remove it from the donor system. Use copies of the completed worksheets as labels:

1. Make a copy of the Worksheets from 2.1 through 2.8 that you filled out for the donor system packet.
   If IOAM enclosures are involved, make a copy of Worksheet 2.1a.
2. To each enclosure you are removing from the donor system, tape a copy of the worksheet that describes that enclosure.

3. Prepare Both ServerNet Fabrics

To successfully reduce your system online, check that both ServerNet fabrics in your system are fully operational. If any ServerNet alarms have been generated, you must eliminate these problems before you can resize your system online.

3.1 Determine Status of ServerNet Fabrics

Use the OSM Service Connection or TSM Service Application and SCF to determine this status, as described in Determine the ServerNet Fabric Status on page A-4.

3.2 Eliminate Problems in the ServerNet Fabrics

If any ServerNet fabrics are not up, follow the recovery procedure in ServerNet Fabrics Are Down on page 7-11.

4. Inventory Enclosures to Be Removed

Before you remove enclosures from your system, inventory each enclosure to identify its devices, communication paths, and subsystems. Record this information on Worksheet 8:

Step	Action	Result
1.	In the tree pane in the OSM Service Connection or TSM Service Application, expand the icon of the highest-numbered enclosure you will remove.	The tree pane displays the devices in the enclosure.
2.	In the tree pane, fully expand the enclosure. For IOAM enclosures, fully expand each module.	The tree pane displays all names of all devices in the enclosure.
3.	Referring to the slot numbers on Worksheet 8, determine the names and types of the associated devices, as follows: Right-click each device listed in the tree pane and fill in the name and type (for example disk, SEB, E4SA) of each device on Worksheet 8. Do the same for each IOAM enclosure, except use Worksheet 8a. Click **Close**.	
4.	Bring Worksheet 8 to the service side of the enclosure. Check the SCSI connectors on the PMF CRUs or IOMF CRUs in slots 50 and 55. If any tape drives are attached to these connectors, add that information to the worksheet.	
5.	Repeat these steps for all other enclosures you plan to remove. Use a separate copy of the worksheet for each enclosure.	

You now can see what devices exist in each enclosure. You can decide:

- Whether to terminate or relocate system resources and user processes
- Which devices to reconfigure on the donor system
- Which devices to move to the target system

For example, if you are removing an enclosure that contains an E4SA, GESA, or G4SA, you know that you might need to move or terminate all TCP/IP connections and any processes started from these sessions.

5. Prepare and Stop Devices and Processes

Because you must maintain access to the system after you have stopped or moved all devices, processes, and their communication paths, the TACL process you use for the following steps should run in group 01. For this reason, use the startup TACL process for system reduction.

<table>
<tr><th>Step</th><th>Action</th></tr>
<tr><td>1.</td><td>Using OSM or TSM and Worksheet 8, collect information about the following devices in the enclosure to be removed:

Internal disk drives
Adapters (in direct or indirect communication with devices):
ServerNet adapters
SEBs or MSEBs (This information is used only for determining communication paths. SEBs and MSEBs are not moved or terminated.)
SWAN concentrators
External tape drives
Open SCSI devices

For IOAM enclosures, collect information about ServerNet Switch Boards, FCSAs, and G4SAs on Worksheet 8a.</td></tr>
<tr><td>2.</td><td>Identify any dependent processes and indirect communication paths before removing an enclosure:
<ol type="a">
To identify currently running processes in a processor enclosure, use the TACL STATUS command on that enclosure. For example:
<code>> STATUS 8</code>
Correlate the list of processes returned with the results of Step 4. Inventory Enclosures to Be Removed on page 4-7.
Identify processes and paths that might be affected.
</td></tr>
<tr><td>3.</td><td>Terminate or move processes as necessary.</td></tr>
</table>

For Information About	See
Establishing a startup TACL connection	Start a Startup TACL Session on page A-9
Examples of preparing and stopping devices and their associated processes	Appendix D, Stopping Devices and Processes
Documentation you might need for your subsystems	Table D-1, References and Examples, on page D-1
Indirect communication paths	Is the Enclosure Providing Critical Resources? on page 3-4 Determining Communication Paths on page 3-6

6. Ensure Devices and Processes Are Stopped

Step	Action	Result
1.	In the tree pane of the OSM Service Connection or TSM Service Application, expand the enclosure and the module of the enclosure that contains the stopped devices, including internal disks.	The tree pane displays the devices in the enclosure.
2.	Make sure all icons are fully expanded.	The tree pane displays all devices in the enclosure.
3.	Verify that the devices you have stopped are in a yellow state.	

For system processes and subsystems, see the appropriate manuals for the commands to show that processes are stopped.

For user processes, use the appropriate tool to verify that processes are stopped.

7. Delete Devices and Processes If Necessary

Skip this step if you will either move these devices and processes to another system or remove them from your system only temporarily.

If you will …	Then …
Remove the enclosure from your system only temporarily	Skip this step. Go to 8. Prepare Enclosures for Removal on page 4-11.
Either: • Never use these devices on the donor system again. • Never add another enclosure later to the donor system that has the same group number as the enclosure being removed.	Perform this step.

7.1 Delete Devices and Processes From System Configuration Database

The commands for deleting devices and processes from the system configuration database are subsystem-specific.

For Information About	See
Documentation to stop devices and processes	Table D-1, References and Examples, on page D-1

7.2 Ensure That Devices Are Deleted

From the OSM Service Connection or TSM Service Application, confirm that all devices have been deleted from the enclosure:

- In the tree pane, all deleted internal disks should be listed as Unknown Disk CRU.
- The Attributes tab should show other devices as Not Configured.

8. Prepare Enclosures for Removal

8.1 Disable Batteries in the Enclosure

This procedure prepares an enclosure for power off by disabling its battery packs.

Step	Action	Result
1.	In the tree pane of OSM or TSM, expand the group and module of the enclosure you want to remove so you can see the batteries.	
2.	Right click the battery in slot 23 and select **Actions**.	The actions dialog box appears.
3.	In OSM Select Enable/Disable. In TSM select Disable Battery. Click **Perform Action.**	
4.	Follow the screen prompts to disable the battery.	
5.	Repeat steps 2 though 4 for the battery in slot 28.	
6.	Close the Actions dialog box	The tree pane reappears.

8.2 Halt Both Processors in a Processore Enclosure

If you are removing a processor enclosure, use this procedure to halt the processors. While performing this procedure, monitor system messages and resolve events if they continue to be repeated after the system has reached a state of low activity.

Step	Action	Result
1.	In the tree pane of OSM or TSM, expand the group and module of the enclosure you want to remove so you can see the PMF CRUs.	
2.	Expand the PMF CRU in slot 50 so you can see the processor.	
3.	Right click the processor and select **Actions**.	The actions dialog box appears.
4.	Select Halt and click **Perform Action**	
5.	Follow the screen prompts to complete the action.	
6.	Repeat steps 2 though 5 for the processor in the PMF CRU in slot 55.	
7.	Close the Actions dialog box	The tree pane reappears.
8.	If not all desired components are halted, the system might not be properly prepared for the rest of the reduction procedure. Before proceeding, consult your system analyst.	

After you perform this procedure for all enclosures to be removed, continue to monitor system messages to determine whether the rest of your system is operating normally.

8.3 Disable ServerNet Ports in Y-fabric MSEBs (or SEBs)

This procedure disables the Y-fabric ServerNet ports in processor enclosures that are connected to the enclosure you are removing.

Note. Consult the topology diagrams in the planning worksheets to determine which ServerNet ports need to be disabled. An I/O enclosure is connected to only one port on one processor enclosure; however, a processor enclosure might be connected to a port on as many as 4 other processor enclosures.

Step	Action	Result
1.	In the tree pane of OSM or TSM, navigate to a processor enclosure that is connected to the enclosure you want to remove. Expand the group and module so you can see the MSEBs (or SEBs).	
2.	Right click the Y-fabric MSEB that is connected to the enclosure you are removing. Select **Actions**.	The actions dialog box appears.
3.	Select Disable ServerNet (Router) Port and click **Perform Action.**	
4.	Follow the screen prompts to disable the appropriate ServerNet port.	
5.	If you are removing a processor enclosure that is connected to other processor enclosures, repeat steps 1 through 4 for the appropriate ServerNet ports on those enclosures.	
6.	Close the Actions dialog box	The tree pane reappears.

Table 4-1. Router Port Connections to SEBs

Router Port	Connects to SEB
1	Cable connector 5
2	Cable connector 4
3	Cable connector 3
4	Cable connector 2
5	Cable connector 1

Table 4-2. Router Port Connections to MSEBs

Router Port	Connects to MSEB
0	Cable connector 6
1	Cable connector 5
2	Cable connector 4
3	Cable connector 3
4	Cable connector 2
5	Cable connector 1

8.4 Remove Cables from Disabled ServerNet Ports on Y-fabric

Looking at an enclosure you are removing, you can see where all the ServerNet cables attach. Therefore, removing that end of the ServerNet cables is simple. However, those cables might be attached at the other end to enclosures that are still operating. You must be careful which ServerNet cables you remove from an enclosure that is still operating. If you disconnect the wrong cable, you might bring down one of the ServerNet fabrics on your system. Therefore, remove only those ServerNet cables that are attached to an enclosure you are removing, as follows:

1. Detach both ends of the ServerNet Y-fabric cables that run between an enclosure you have prepared and the rest of the system.
2. Using the OSM Service Connection or TSM Service Application and SCF, check the status of the ServerNet Y fabric, as described in Determine the ServerNet Fabric Status on page A-4.

   It is normal for a ServerNet alarm to be generated on the link you have just removed. However, if an alarm is generated on any other links, you have detached a live ServerNet cable: one that does not connect on at least one end to an enclosure that you powered off.

   Repair the problem that generated the ServerNet alarm. Because problems in ServerNet fabrics are often specific to an individual system, this guide cannot contain every diagnostic or recovery procedure. Consult your service provider.

8.5 Power Off Enclosure

Ensure that all processors in the enclosures to be removed have been prepared and halted, if applicable, before powering off the enclosures:

Step	Action	Result
1.	From the OSM Service Connection or TSM Service Application, select the Display menu. Right-click the enclosure you want to power off.	A dialog box appears.
2.	Select **Actions**.	A dialog box appears.
3.	Select **Power Off**.	
4.	Select **Perform Action**.	
5.	Click **OK** and then **Close**. In OSM, if the battery is not disabled, another confirmation dialog box opens. If this happens, click **Yes**.	In the tree pane, the devices in the enclosure to be removed disappear. The icon of the enclosure disappears from the Management window, indicating that the enclosure has been powered off.
6.	Repeat this procedure for each enclosure to be removed.	

△ **Caution.** For TSM, if an enclosure does not power off when you perform this procedure, you can perform a Force Power Off operation. However, the Force Power Off operation might have severe consequences. Before you perform this operation, consult your system analyst. For OSM, if an enclosure does not power off when you perform this procedure, contact your service provider.

Note. You cannot power off an IOAM enclosure from OSM. Under normal circumstances, only a service provider trained by HP should power off an IOAM enclosure.

8.6 Disable ServerNet Ports in X-fabric MSEBS

This procedure disables the X-fabric ServerNet ports in processor enclosures that are connected to the enclosure you are removing.

Note. Consult the topology diagrams in the planning worksheets to determine which ServerNet ports need to be disabled. An I/O enclosure is connected to only one port on one processor enclosure; however, a processor enclosure might be connected to a port on as many as 4 other processor enclosures.

Step	Action	Result
1.	In the tree pane of OSM or TSM, navigate to a processor enclosure that is connected to the enclosure you want to remove. Expand the group and module so you can see the MSEBs (or SEBs).	
2.	Right click the X-fabric MSEB that is connected to the enclosure you are removing. Select **Actions**.	The actions dialog box appears.
3.	Select Disable ServerNet Port and click **Perform Action.**	
4.	Follow the screen prompts to disable the appropriate ServerNet port.	
5.	If you are removing a processor enclosure that is connected to other processor enclosures, repeat steps 1 through 4 for the appropriate ServerNet ports on those enclosures.	
6.	Close the Actions dialog box	The tree pane reappears.

Table 4-3. Router Port Connections to SEBs

Router Port	Connects toSEB
1	Cable connector 5
2	Cable connector 4
3	Cable connector 3
4	Cable connector 2
5	Cable connector 1

Table 4-4. Router Port Connections to MSEBs

Router Port	Connects to MSEB
0	Cable connector 6
1	Cable connector 5
2	Cable connector 4
3	Cable connector 3
4	Cable connector 2
5	Cable connector 1

8.7 Remove Cables from Disabled ServerNet Ports on X-fabric

Looking at an enclosure you are removing, you can see where all the ServerNet cables attach. Therefore, removing that end of the ServerNet cables is simple. However, those cables might be attached at the other end to enclosures that are still operating. You must be careful which ServerNet cables you remove from an enclosure that is still operating. If you disconnect the wrong cable, you might bring down one of the ServerNet fabrics on your system. Therefore, remove only those ServerNet cables that are attached to an enclosure you have are removing, as follows:

1. Detach both ends of the ServerNet X-fabric cables that run between an enclosure you have prepared and the rest of the system.
2. Using the OSM Service Connection or TSM Service Application and SCF, check the status of the ServerNet X fabric, as described in Determine the ServerNet Fabric Status on page A-4.

   You should not normally see alarms at this point. However, if a ServerNet alarm is generated, you might have detached a live ServerNet cable: one that does not connect on at least one end to an enclosure that you powered off.

   Repair the problem that generated the ServerNet alarm. Because problems in ServerNet fabrics are often specific to an individual system, this guide cannot contain every diagnostic or recovery procedure. Consult your service provider.

8.8 Reenable Disabled ServerNet Ports on both fabrics

This procedure enables the ServerNet ports that were disabled in steps 8.3 and 8.6.

Note. Consult the topology diagrams in the planning worksheets to determine which ServerNet ports need to be reenabled. If you are unsure whether a port needs to be to enabled, enable it anyway. Enabling an already-enabled port does not cause any problems.

Step	Action	Result
1.	In the tree pane of OSM or TSM, navigate to the processor enclosure to which the removed enclosure was connected. Expand the group and module so you can see the MSEBs (or SEBs).	
2.	Right click the Y-fabric MSEB and select **Actions**.	The actions dialog box appears.
3.	Select Enable ServerNet (Router) Port and click **Perform Action.**	
4.	Follow the screen prompts to disable the appropriate ServerNet port.	
5.	Right click the X-fabric MSEB and select **Actions**.	The actions dialog box appears.
6.	Select Enable ServerNet (Router) Port and click **Perform Action.**	
7.	Follow the screen prompts to disable the appropriate ServerNet port.	
8.	If you are removing a processor enclosure that is connected to other processor enclosures, repeat steps 1 through 7 for the appropriate ServerNet ports on those enclosures.	
9.	Close the Actions dialog box	The tree pane reappears.

8.9 Verify Enclosures Are Removed From System Configuration

The icons of the enclosures disappear from the display in the OSM Service Connection or TSM Service Application, but you need to confirm that the enclosures have been removed from the system configuration:

Step	Action	Result
1.	Log off of the OSM or TSM Low-Level Link Application and then log back on.	This action synchronizes the OSM Service Connection or TSM Service Application with the resized system.
2.	From the OSM or TSM Low-Level Link Application, click **System Discovery.**	The Management window appears. The display might take several minutes to appear.
3.	In the tree pane, double-click the icon for the system.	All enclosures appear except the removed enclosures. If all remaining enclosures with their devices do not appear after several minutes, see System Discovery Does Not Show All CRUs or Enclosures on page 7-13.

8.10 Prepare Any Remaining Enclosures To Be Removed

Repeat steps 8.1 through 8.9 for all enclosures you are removing.

9. Finish the Reduction

9.1 Save New Donor System Configuration

Now that you have altered the donor system's configuration, save its current configuration.

1. At a TACL prompt, start SCF and enter the SAVE CONFIGURATION command:

   ```
   >SCF
   -> SAVE CONFIGURATION xx.yy
   ```

   where *xx.yy* is a configuration number in the range 00.00 through 99.99 that is both:

   - Not already in use
   - Not the configuration number you used in Step 1.3 Save the Current Donor System Configuration on page 4-3

   If you specify a number that is already in use, you receive this message:

   ```
   Configuration file $SYSTEM.ZSYSCONF.CONFxx.yy already exists.
   Replace it (Y/[N])
   ```

2. Enter *N* and retry the SAVE CONFIGURATION *xx.yy* command using another number.

3. Record the number of your saved configuration (*xx.yy*) on Worksheet 1.

9.2 Maintain System Console Connections

If your donor system will remain in use, it is recommended that you maintain the following sessions on a system console so that you can monitor events on this donor system:

- CLCI (Startup TACL) connection
- CNSL (EMS Event Stream) connection
- OSM or TSM Low-Level Link
- OSM Service Connection or TSM Service Application

10. Remove Other Cables From Powered-Off Enclosures

You must remove the remaining cables between the donor system and the powered-off enclosures.

For Information About	See
The other cables you need to remove	Cabling on page 1-8
Procedures for removing cables	*NonStop S-Series Hardware Installation and FastPath Guide*
Cables for an IOAM enclosure	*Modular I/O Installation and Configuration Guide*

11. Physically Remove Enclosures From the System

△ **Caution.** Enclosures are not customer-replaceable units (CRUs). Removing enclosures should be done only by service providers trained by HP.

For Information About	See
Procedures for removing enclosures	*NonStop S-Series Service Provider Supplement* (available to service providers)

Adding Enclosures to Another System

If you will never use these enclosures again on any system, you are finished with this guide.

For Information About	See
Adding these enclosures online to another systems	Section 6, Expanding a System Online
Adding these enclosures offline to another system	*NonStop S-Series Planning and Configuration Guide*

5 Reducing a System Offline

This section describes the procedure for removing enclosures offline.

Reducing a system offline is removing enclosures from a donor system after you have shut down the entire system. You can do this on any RVU. For systems running an RVU earlier than G06.04, this is the only way to reduce a system.

The procedure for reducing your system offline is:

Step	As Described In	Page
1. Gather all information and prepare your system.	1. Prepare the Donor System for Reduction	4-2
	2. Record Information About the Donor System	4-4
	4. Inventory Enclosures to Be Removed	4-7
	5. Prepare and Stop Devices and Processes	4-8
	6. Ensure Devices and Processes Are Stopped	4-9
	7. Delete Devices and Processes If Necessary	4-9
2. Power off your system by powering off each enclosure.	8.9 Verify Enclosures Are Removed From System Configuration	4-18
3. Remove enclosures.	*NonStop S-Series Service Provider Supplement*	
4. Use the OSM or TSM Low-Level Link to load the reduced donor system.		

Note. When you power your system on again, you might need to reset the system time. For the procedure, see the *SCF Reference Manual for the Kernel Subsystem.*

When you finish the offline removal of enclosures from the donor system, you are ready to expand the target system. Proceed to Section 6, Expanding a System Online.

6 Expanding a System Online

This section describes the procedure for adding enclosures online to expand a system.

You must perform each step as documented in the order shown. If you do not, system expansion could fail with unpredictable results.

Preparation for Online Expansion

Preparation for online expansion depends on where you get the enclosures with which you will be expanding your system.

If the enclosures come from another system:

Step	For more information
1. Ensure that the system you are expanding contains an adequate number of MSEBs and MSEB PICs.	Does Your System Have Enough MSEBs and MSEB PICs? on page 1-3
2. Plan the procedure of removing enclosures from another system.	Section 3, Planning System Reduction
3. Remove enclosures from the donor system.	Section 4, Reducing a System Online
4. Expand the target system.	Perform the procedures in this section.

If the enclosures come directly from the factory:

Step	For more information
1. Ensure that the system you are expanding contains an adequate number of MSEBs and MSEB PICs.	Does Your System Have Enough MSEBs and MSEB PICs? on page 1-3
2. Expand your system.	Perform the procedures in this section.

△ **Cautions.**

- You can expand your system only if it already contains an adequate number of MSEBs and MSEB PICs to support the increased number of processors after expansion.
- You can expand your system online only if it is running G06.00 or later with a compatible TSM package or G06.08 or later with a compatible OSM or TSM package.

  For information about the OSM package and other requirements for migrating to OSM, see the *OSM Migration Guide*.

  For information about the TSM package, see the TSM Read Me file at **Start > Programs >Compaq TSM > TSM Documentation > TSM Read Me > TSM Product Version Information > TSM Software Product Versions**. If you are running an RVU earlier than G06.00, you must expand the system offline, as described in the *NonStop S-Series Planning and Configuration Guide*.
- If you are expanding your system to include IOAM enclosures, you need the G06.24 RVU or later and the OSM software components that are included with it. The IOAM enclosure and its components must be installed by a service provider trained by HP.
- If you are expanding your system, you might need a different topology. **Topologies cannot be changed online**. After you configure your system for a particular topology, changing the topology requires stopping your system. For more information, see the *NonStop S-Series Planning and Configuration Guide*.
- It is critical to set the topology value (Tetra 8 or Tetra 16) correctly on each enclosure you add. These values must match the value on the target system.
- Ensure you have KMSF swap files of the proper size and location available on both systems. Otherwise, applications and processes on the newly resized systems might terminate. For information about KMSF swap files and their recommended sizes, see the *Kernel-Managed Swap Facility (KMSF) Manual*.
- When you attach ServerNet cables between the block and the target system, you must attach all cables correctly the first time. If you attach the cables to the wrong ports, perform the recovery described in ServerNet Fabrics Are Cross-Cabled on page 7-10.
- The SP firmware on all blocks to be added must be identical to the SP firmware on your target system.
- Do not unplug and replug any AC power cords. If you do, see AC Power Cord Was Unplugged or Replugged on page 7-2. After any enclosure has been plugged into an AC power source once, unplugging and then replugging its AC power cord powers on the enclosure prematurely.
- If you are expanding your system with enclosures taken from another system, plug in their AC power cords last, after you have completed all other steps.

1. Prepare Target System for Expansion

This procedure assumes that both ServerNet fabrics in your system are fully operational. If any ServerNet alarms have been generated, you must eliminate those problems before you expand your system online.

Adding blocks must not conflict with the operations of your system. Therefore, you must prepare both the blocks you are adding and the target system you are expanding.

For...	See...
Tools and Supplies	Table 6-1, following
Preparatory Steps	Table 6-2 on page 6-5
Checklists to copy and mark	Appendix C, Checklists and Worksheets

Table 6-1. Tools and Supplies

Tool or Supply	Used to …
A system console, as described in Step 5.1 Choose System Console to Use on page 6-11	Communicate with target system and with the enclosures you are configuring into a block
HP NonStop Server System Console Installer CD	Install the OSM or TSM package on all system consoles
Ethernet switch or hub and three Ethernet cables	Connect the system console to the enclosures you are configuring into a block
One or more NonStop S-series enclosures	Expand your system
SEBs or MSEBs	Expand your system
Adapters or other components	Populate added blocks on your system as needed
Short Phillips screwdriver	Tighten screws on power retainers and groundstraps
Small slotted screwdriver	Tighten screws
3/16-inch hex socket driver	Tighten nuts
3/4-inch or 9/16-inch open-end wrench	Lower the enclosure legs
Antistatic mat	Protect components from electrostatic discharge
Tape or cable labels	Label ServerNet cables and attach worksheets to enclosures
Scissors or cutter	Cut plastic straps on shipping packages

Table 6-2. Preparatory Steps

Step	Action	Reason
1.	If necessary, upgrade the OSM or TSM client software to a compatible version on all system consoles used in the resizing process.	For online resizing to succeed, every system console involved must run identical versions of the OSM or TSM package. To determine which product versions of the OSM package are compatible with your system, see the *OSM Migration Guide.* To determine which product versions of the TSM package are compatible with your system, see the TSM Read Me file at **Start > Programs > Compaq TSM > TSM Documentation.** For instructions about upgrading the OSM or TSM client software, see the *NonStop System Console Installer Guide*.
2.	If you are expanding your system with enclosures from another system, have handy the donor system packet, as described in Step 1.1 Prepare Worksheets on page 4-2.	You need a record of configuration information about the enclosures from the donor system.
3.	If you are planning a block that contains more than three enclosures, make an extra copy of Worksheet 7 for every two additional enclosures.	Worksheet 7 records information for a maximum of one processor enclosure and two I/O enclosures. Larger blocks require more copies of this worksheet.
4.	Place the enclosures to be added near the target system in the same physical arrangement they will have after resizing is complete.	This practice minimizes recabling and troubleshooting.
5.	Label all ServerNet cables.	Do not cross-cable the ServerNet X and Y fabrics. Figure 6-1 shows an example of a simple label: a piece of tape with *X* or *Y* written on it, attached to the cable.

Figure 6-1. Examples of Labeled ServerNet Cables

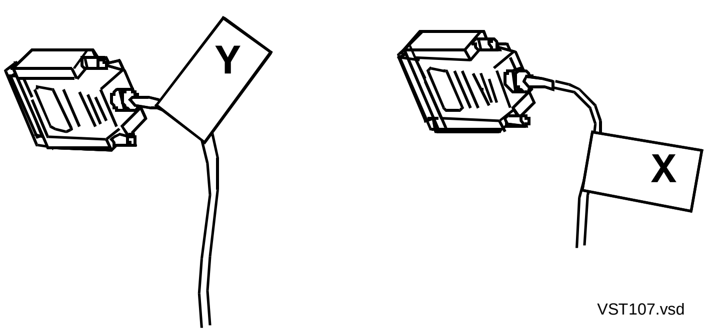

2. Record Information About Target System

Before you physically add enclosures to your target system, you must record information about the target system on worksheets in Appendix C. You determine this information with a system console.

For information about …	See …
Starting the OSM or TSM Low-Level Link	Start the OSM or TSM Low-Level Link on page A-9
Locating OSM or TSM views	OSM or TSM online help

While you are recording information for the target system, someone else can be recording the same information for the enclosures to be added to the target system. This process is described in Step 6. Finish Gathering Information on page 6-14.

2.1 Prepare Worksheets

1. From Appendix C, make copies of Worksheet 1, Worksheets 2.1 through 2.8, and Worksheet 7. If the target system contains any IOAM enclosures, also copy Worksheet 2.1a.
2. Circle the words *Target System* at the top of each worksheet.
3. Staple these worksheets together into a **target system** packet.

Keep this packet separate from the worksheets for the donor system, if any. Do not mix information from your target system with information from your donor system.

2.2 Record Information About the Master Service Processor

Record the following information for the master service processor in group 01:

Step	Action	Result
1.	In the OSM or TSM Low-Level Link window, click **System Discovery**.	After several minutes, the Management window appears. The Attributes tab displays the following: ● System Name ● System Serial Number ● System Type ● Topology ● SP Firmware Version (on the status bar at the bottom of the main window)
2.	Record this information on Worksheets 1 and 7 of the target system packet.	

△ **Caution.** If an enclosure in the target system has a different topology from that of group 01, this information-gathering procedure will not find the incompatibility. However, this incompatibility can prevent successful expansion. Recovery requires a system outage and resynchronizing the topology of all enclosures in the target system.

2.3 Record Information About Other Processors

Perform this step twice for each enclosure in the target system:

- Once for the PMF CRU or IOMF CRU in slot 50
- Once for the PMF CRU or IOMF CRU in slot 55
- For an IOAM enclosure, once for each ServerNet Switch Board

Step	Action	Result
1.	In the tree pane, located on the left side of the window, expand the icon for the system.	The tree pane displays all the enclosures.
2.	Expand the group and module so that you can see the PMF or IOMF CRUs. For an IOAM enclosure, expand the group, the IOAM enclosure, and the modules so that you can see the ServerNet Switch Boards.	The Physical view displays the devices in the enclosure.
3.	Expand the PMF or IOMF in slot 50 and click on the SP. For an IOAM Enclosure, click on the ServerNet Switch Board in module 2.	The attributes appear in the detail pane, including: • IP Address • SP Firmware Version If this is a processor enclosure, you also see: • Processor Firmware Version • SCSI Firmware Version • The Topology is either 8 (for Tetra 8) or 16 (for Tetra 16). If this is a ServerNet Switch board, you see: • ME Firmware version • ME FPGA version
4.	Record this information on Worksheet 1 and 2.*n* of the target system packet. For IOAM enclosures, use Worksheet 2.1a of the target system packet.	
5.	For the CRU in slot 55, repeat steps 3 and 4. For the ServerNet Switch Board in module 3 of an IOAM enclosure, repeat steps 3 and 4.	
6.	For every enclosure in the target system, repeat steps 2 through 5.	

Note. IP addresses for the ServerNet Switch boards in an IOAM enclosure are dynamically assigned by a DHCP server or software that performs the same function.

2.4 Record Service Connection IP Addresses

Perform this step once for each enclosure in the target system.

Step	Action	Result
1.	In the OSM or TSM Low-Level Link window, select **File.**	A menu appears.
2.	Select **Edit System List ….**	The Edit Systems List dialog box appears, containing the name of the target system. (You might have to scroll down to find it.)
3.	Highlight the name of the target system.	
4.	Click **Modify.**	The Modify System dialog box displays the following information: ● Primary Service Connection IP address ● Backup Service Connection IP address
5.	Record these addresses on Worksheet 2.*n* of the target system packet.	
6.	Click **Cancel**.	The Modify System dialog box disappears.
7.	Click **OK.**	The Edit Systems List dialog box disappears.
8.	Repeat this procedure from Step 1 for every other enclosure in the target system.	

3. Prepare Target System for Addition of Block

On the target system, start a TACL session and log on as a super-group user.

3.1 Record Current SYSnn Subvolume

1. At the TACL prompt:

   ```
   > STATUS *, TERM, PROG
   ```

2. Record the SYS*nn* subvolume on Worksheet 1.

3.2 Check Storage Automatic Configuration and Autostart

It is recommended that you start an SCF log file to track all configuration changes made during the system expansion.

1. Start SCF and enter the following command:

   ```
   -> INFO SUBSYS $ZZSTO
   ```

   You will see something similar to the following display:

   ```
   AutoConfigure  AutoRevive  AutoStart  BulkIO   LabelTape  UPS
   ON             OFF         OFF        OFF      ON         OFF
   ```

   In this display:

   - AutoConfigure is ON, which indicates that the Storage Automatic Configuration feature of this system is on.
   - AutoStart is OFF, which indicates that the Autostart feature of this system is off.

   If you set both the Storage Automatic Configuration and Autostart features of your target system to ON, all internal disk drives in that block are automatically configured and started when you add the block to your system.

2. If Step 1 shows that either AutoConfigure or AutoStart on the target system is OFF, enable these features by entering the following SCF command:

   ```
   -> ALTER SUBSYS $ZZSTO, AUTOCONFIG ON, AUTOSTART ON
   ```

   For additional information about the Storage Automatic Configuration and Autostart features of internal disks, see the *SCF Reference Manual for the Storage Subsystem*.

4. Save Current Target System Configuration

Before you alter your target system's configuration, save its current configuration for fallback purposes:

1. At the SCF prompt, enter the SAVE CONFIGURATION command:

   ```
   -> SAVE CONFIGURATION xx.yy
   ```

   where *xx.yy* is a configuration number in the range 00.00 through 99.99 that is not already in use.

   If you specify a number that is already in use, you receive this message:

   ```
   Configuration file $SYSTEM.ZSYSCONF.CONFxx.yy already exists.
   Replace it (Y/[N])
   ```

2. Enter *N* and retry the SAVE CONFIGURATION *xx.yy* command with another number.

3. Record the number of your saved configuration (*xx.yy*) on Worksheet 1.

5. Copy SP Firmware File From the Target System to the System Console

In this step, you shift your activities from the target system to the block you will add. If you have a system console to move from the target system to that block, this step indicates when to move the system console.

The block you add to your target system must have SP firmware running on the dedicated LAN that is identical to the SP firmware on the target system. This step copies the target system's SP firmware file onto a system console. This file, in turn, is copied to all enclosures in the block.

5.1 Choose System Console to Use

You can use one of the following system consoles:

- One of the two system consoles attached to your donor or target system.
- The one system console that is attached to your target or donor system. However, if you use the only system console from either system for this operation, you leave that system without dial-in/dial-out capability until you reconnect the system console.

When you have chosen the system console you will use, attach it to the target system if necessary, using the procedure Move the System Console on page A-8.

5.2 Identify IP Address for FTP Session

You will copy the SP firmware file from the target system to the system console through a dedicated communications adapter such as an E4SA. On the system console on the target system, determine the IP address of the adapter you will use:

1. Log on to TACL and start SCF.
2. Enter the following SCF command to display IP addresses for the dedicated communications adapters on the target system:

```
-> INFO SUBNET $*.*
```

A display such as the following can take a few moments to appear:

```
->INFO SUBNET $*.*

TCPIP Info SUBNET \TROLL.$ZTC00.*

Name     Devicename     *IPADDRESS       TYPE     *SUBNETMASK  SuName   QIO *R

#LOOP0   \NOSYS.$NOIOP   127.0.0.1       LOOP-BACK %HFF000000           OFF N
#SN1     \TROLL.LAN670A  172.17.116.175  ETHERNET  %HFFFFFF00           ON  N

TCPIP Info SUBNET \TROLL.$ZTC20.*

Name     Devicename     *IPADDRESS       TYPE     *SUBNETMASK  SuName   QIO *R

#LOOP0   \NOSYS.$NOIOP   127.0.0.1       LOOP-BACK %HFF000000           OFF N
#SN1     \TROLL.LAN020A  172.17.116.190  ETHERNET  %HFFFFFF00           ON  N

TCPIP Info SUBNET \TROLL.$ZTCP0.*

Name     Devicename     *IPADDRESS       TYPE     *SUBNETMASK  SuName   QIO *R

#LOOP0   \NOSYS.$NOIOP   127.0.0.1       LOOP-BACK %HFF000000           OFF N
#SN1     \TROLL.LANX     172.17.116.173  ETHERNET  %HFFFFFF00           ON  N

TCPIP Info SUBNET \TROLL.$ZTCP1.*

Name     Devicename     *IPADDRESS       TYPE     *SUBNETMASK  SuName   QIO *R

#LOOP0   \NOSYS.$NOIOP   127.0.0.1       LOOP-BACK %HFF000000           OFF N
#SN1     \TROLL.LANY     172.17.116.174  ETHERNET  %HFFFFFF00           ON  N
```

3. Note the adapters named *system-name*.LANX or *system-name*.LANY. Do not use either of these adapters to copy SP firmware.

4. Ask your network administrator which one of the other adapters should be used to create a file transfer protocol (FTP) session between the system console and the target system. In this example, the E4SA device is \TROLL.LAN670A, shown in boldface. Its IP address of 172.17.116.175 is used in the next step.

5.3 Copy SP Firmware File From Target System

On the system console attached to the target system:

1. Select **Start > Programs > Command Prompt**
2. Create a directory for the SP firmware file on the target system console:

   ```
   C:\>md \spcode9
   ```
3. Switch to that directory:

   ```
   C:\>cd \spcode9
   ```
4. Start an FTP session
   a. **Start > Programs > Command Prompt**
   b. Enter **ftp** and the IP address obtained from Step 5.2 Identify IP Address for FTP Session on page 6-11. For example:

      ```
      C:\>ftp 172.17.116.175
      ```
   c. Log on as a super-group user:

      ```
      User (172.17.116.175:(none)): super.compaq
      Password:
      ```
   d. Enable the Guardian API:

      ```
      ftp> quote guardian
      ```
5. Move to the current SYS*nn* subvolume of your target system. To be sure of the SYS*nn* you want, see Worksheet 1.

   ```
   ftp>cd $SYSTEM.SYSnn
   ```
6. Prepare for a file transfer:

   ```
   ftp>bin
   ```
7. Retrieve the SP firmware file to your current directory:

   ```
   ftp>get spntcde9
   ```
8. Log off of the FTP session:

   ```
   ftp>bye
   ```

6. Finish Gathering Information

At this stage, you treat and configure each enclosure separately. Before you configure an enclosure into a block, you must supply its IP information. IP information for an enclosure is recorded on different worksheets, depending on whether the enclosure came directly from the factory or from a donor system.

Note. IP addresses for the ServerNet switch boards of the IOAM enclosures are dynamically assigned by a DHCP server or software that performs the same function.

Make a copy of Worksheet 7 and keep it accessible. You will use it later when you configure the enclosures into a block.

If Enclosure Comes From the Factory

If the enclosure was shipped directly from the factory, it contains default IP addresses. These addresses are listed in Table 6-3 and preprinted on Worksheet 7.

Table 6-3. Default MSP IP Addresses

Item	Value
Slot 50	192.231.36.2
Slot 55	192.231.36.3
Subnet Mask	255.255.255.0
Gateway Route	192.231.36.9

If Enclosure Comes From a Donor System

If the enclosure is coming from a donor system, you already recorded the slot 50 and slot 55 IP addresses on Worksheet 2.*n* of the donor system packet during Step 2.1 Record Basic Information on page 4-4.

Contact your network administrator to identify the subnet mask and gateway route IP addresses for the donor system.

On Worksheet 7 of the donor system packet, record these values and the information from Worksheet 1.

7. Connect a System Console to the Enclosure

Caution. Before you perform this step, ensure that:

- The enclosure is powered off.
- The system console that you will connect to the enclosure is running a version of the OSM or TSM client software that is compatible with your RVU and TSM server software.

  To determine the product version of OSM or TSM client software on the system console, see Determine the Product Version of the TSM Client Software on page A-7.

  For instructions about how to upgrade the OSM or TSM client software on the system console, see the *NonStop System Console Installer Guide* for the instructions for the version you plan to install.

Both the system console that you will connect to the enclosure and the enclosure itself must be offline (and therefore isolated). Until you configure the enclosure into a block, it cannot share a network with the target system.

Using the procedure Move the System Console on page A-8, connect to the enclosure the system console selected in Step 5.1 Choose System Console to Use on page 6-11.

Note. To migrate a system console to the Microsoft Windows 2000 Professional operating system, see *NonStop System Console Guide for Migrating to Microsoft Windows XP Professional.* For additional information, see Hotstuff HS02701A.

8. Change Group Number of Enclosure to 01

Before you configure an enclosure into a block, you must temporarily change the enclosure's group number to 01.

Note. To change the group number of IOAM enclosures, your service provider uses the Configure Module action in the OSM Low-Level Link Application.

Change I/O enclosure numbers with the group ID switches. These switches are shown in Figure 6-2.

To change the group ID switches on I/O enclosures, use a flashlight if necessary to locate the holes near the switches, and press a pointed object such as a straightened metal paper clip into the holes. Do not use a pencil. The pencil point could break and jam the switches.

Press one or both of the holes below the group ID switch to increment the I/O enclosure number, or press one or both of the holes above the group ID switch to decrement the I/O enclosure number.

The group IDs on the group ID switches must be identical. If not, you cannot log on to the OSM or TSM Low-Level Link.

Figure 6-2. Group ID Switches

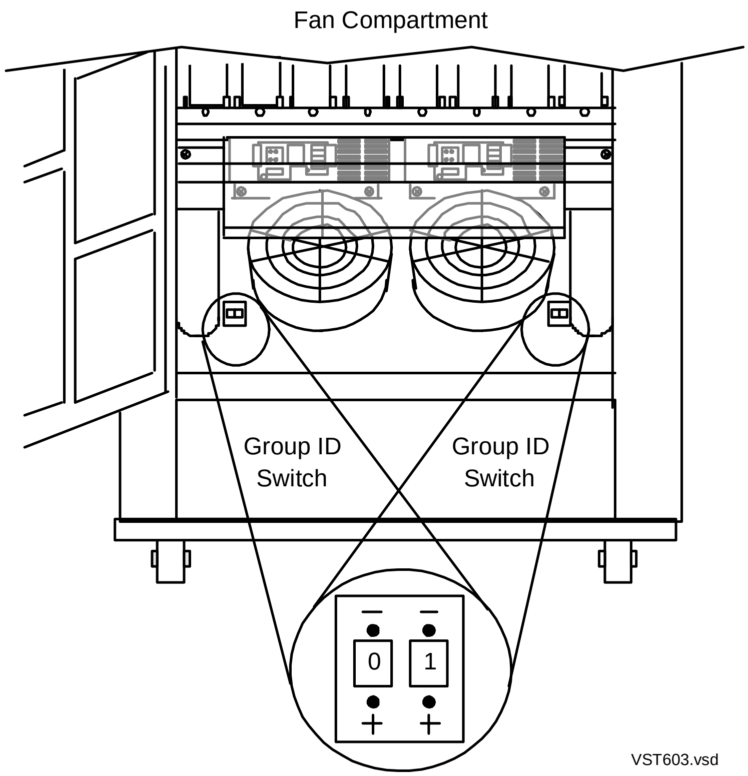

9. Power On Enclosure

Power on this enclosure by plugging the AC power cords for the enclosure into the AC power source.

△ **Caution.** After an enclosure has been plugged into an AC power source once, unplugging and then replugging its AC power cord powers on the enclosure prematurely. Therefore, do not unplug and replug any AC power cords.

If you are expanding your system with enclosures taken from another system, plug in their AC power cords last, after you have completed all other steps.

If you plug in an AC power cord prematurely (or plug it in and then unplug it), see AC Power Cord Was Unplugged or Replugged on page 7-2.

All CRUs in the enclosure should start automatically. They might take several minutes to run their power-on self-tests (POSTs). Verify that all CRUs are operating:

- The amber LEDs on disk drives and on CRUs installed in slots 50 through 55 will become lit briefly as these CRUs are brought up. However, if any of these amber LEDs remain lit for more than about seven to ten minutes, see POST Fails on page 7-9.
- Green LEDs should be lit.
- The fans on the appearance side of the enclosure should be working.

  If ambient noise prevents you from hearing whether the fans are working, check for air flow by holding your hand at the outlet grill located on the service side of the enclosure just above slots 51 through 54.

After plugging in the AC power cords for the enclosure, if the enclosure does not power on, see Added Enclosure Does Not Power On on page 7-3.

After applying power to the enclosure, if the LEDs on a PMF CRU or IOMF CRU remain unlit:

- Check that the power-on cables are properly attached so that each PMF CRU or IOMF CRU can send the power-on signal to the next. For cabling patterns, see the *NonStop S-Series Hardware Installation and FastPath Guide*.
- If all power-on cables are attached correctly but the CRUs still do not start, see CRU Does Not Start on page 7-6.

Note. For problems with components in an IOAM enclosure, your service provider should see the *Modular I/O Installation and Configuration Guide*.

---

10. Verify Connection Between System Console and Enclosure

In this step, you check whether the system console and enclosure you have connected are communicating.

10.1 Verify Connection to ServerNet X Fabric

Open a Command Prompt window on the system console and enter the following command using the low-level link IP address of group 01, slot 50, from Worksheet 2.1:

C:\>ping *IP-address-of-the-SP-in-slot-50*

The following example is the display you see if the system console is connected to the ServerNet X fabric:

```
C:\>ping 192.231.36.2

Pinging 192.231.36.2 with 32 bytes of data:

Reply from 192.231.36.2: bytes = 32 time =  10ms TTL = 60
Reply from 192.231.36.2: bytes = 32 time =  30ms TTL = 60
Reply from 192.231.36.2: bytes = 32 time = <10ms TTL = 60
Reply from 192.231.36.2:   bytes = 32 time =  30ms TTL = 60
```

If you receive any other messages, see Ping Command Does Not Work on page 7-7.

10.2 Verify Connection to ServerNet Y Fabric

Enter the following command with the low-level link IP address of group 01, slot 55, from Worksheet 2.1:

C:\>ping *IP-address-of-the-SP-in-slot-55*

The following example is the display you see if the system console is connected to the ServerNet Y fabric:

```
C:\>ping 192.231.36.3

Pinging 192.231.36.3 with 32 bytes of data:

Reply from 192.231.36.3: bytes = 32 time =  60ms TTL = 60
Reply from 192.231.36.3: bytes = 32 time =  11ms TTL = 60
Reply from 192.231.36.3: bytes = 32 time =  40ms TTL = 60
Reply from 192.231.36.3: bytes = 32 time = <10ms TTL = 60
```

If you receive any other messages, see Ping Command Does Not Work on page 7-7.

11. Configure System Console and Enclosure

This step ensures that the system console attached to this enclosure contains the correct software and that this enclosure contains the correct IP addresses.

11.1 Check and Upgrade OSM and TSM Client Software If Necessary

You must check to see whether the system console attached to this enclosure is running the appropriate OSM or TSM client software. If necessary, you must upgrade this software.

1. To determine the version of OSM or TSM client software on the system console, see Determine the Product Version of the TSM Client Software on page A-7.
2. If you need to upgrade OSM or TSM client software on the system console, see the *NonStop System Console Installer Guide* for the instructions for that version.

11.2 Enter Enclosure IP Addresses

Step	Action	Result
1.	Start the OSM or TSM Low-Level Link, but do not log on.	
2.	When the logon dialog box appears, click **Cancel.**	
3.	From the menu bar, select **File.**	A menu appears.
4.	Select **Edit System List….**	The Edit Systems List dialog box appears.
5.	Click **Add**.	The Add System dialog box appears.
6.	Using Worksheet 2.*n*, enter the low-level link IP addresses of the PMF CRUs or IOMF CRUs in slots 50 (primary) and 55 (backup) of the enclosure. You do not need to fill in the Service Connection IP addresses.	
7.	After you enter the IP addresses, click **Add**.	
8.	If you see a warning message that an IP address already appears in your system, click **Yes.**	You will overwrite this IP address with the identical IP address that is associated with the added enclosure. If you are using only one system console and moving it between the target system and the enclosure, you must add the other IP address again when you reattach the system console to the target system.
9.	Click **OK.**	The Edit Systems List dialog box disappears.

12. Verify SP Firmware Is Compatible

This step checks whether the SP firmware on the enclosure is compatible with the SP firmware on the target system.

Note. This step does not apply to IOAM enclosures.

12.1 Determine SP Firmware Version for Enclosure

How you perform this step depends on where you got the enclosure:

- If this enclosure was shipped directly from the factory:

Step	Action	Result
1.	Log on to the OSM or TSM Low-Level Link, and click **SP Actions**.	The Service Processor Actions dialog box appears. It might take some time to retrieve all the information. In the Firmware Version list, you will see the SP firmware version for the enclosure. You might need to scroll to find this information.
2.	Record this SP firmware version on Worksheets 1 and 7. For example: `T1089G05^03FEB99^19FEB99^AAP`	

- If this enclosure came from a donor system, locate Worksheet 7 and note the SP firmware version that you recorded for this enclosure.

12.2 Check Compatibility of the SP Firmware

You have recorded on Worksheet 7 SP firmware version information for both your target system (from Step 2.2 Record Information About the Master Service Processor on page 6-6) and for this enclosure (Step 12.1 Determine SP Firmware Version for Enclosure on page 6-20).

1. Compare the values for these two versions.
2. Note whether they match.

13. Update SP Firmware in Enclosure If Necessary

Note. For IOAM enclosures, your service provider must verify that the ME firmware and the ME FPGA is the latest version. Updating the ME firmware and ME FPGA is described in the *Modular I/O Installation and Configuration Guide*.

You might be able to skip this step:

If SP Firmware Versions Are …	Then …
Identical on the enclosure to be added and the target system	Skip this step. Go to Step 14. Configure Topology of Enclosure If Necessary on page 6-22.
Different on the enclosure to be added and the target system	Perform this step.

To update the SP firmware on the enclosure to be added, complete the following procedure for both its CRUs:

Step	Action	Result (page 1 of 2)
1.	On the system console attached to this enclosure, from the OSM or TSM Low-Level Link, click **SP Actions**. If system discovery has occurred, select **Display > SP Actions**	The Service Processor Actions dialog box appears.
2.	In the Service Processors (SPs) list, select the service processor in slot 50 for the enclosure you are adding.	
3.	In the Actions list, select **Firmware Update.**	
4.	Click **Browse.**	The Open dialog box appears.
5.	Navigate to and open the **spcode9** directory (created in Step 2 on page 6-13).	
6.	Double-click the **spntcde9** SP firmware file.	
7.	Click **Perform Action.**	A warning appears that the update might take up to 30 minutes.
8.	Click **OK.**	The status line tracks the update.
9.	When the status of the update is Completed, click **OK**.	The Actions menu of the Service Processor Actions dialog box reappears.
10.	Select **Reset.**	
11.	Click **Perform Action.**	
12.	Click **OK.**	Even though the request is labeled Completed, the action itself takes up to two minutes to finish.

Step	Action	Result (page 2 of 2)
13.	Click **Refresh** after approximately one minute.	If the correct version of the SP firmware is displayed in the Service Processor (SPs) list, the update is successful.
14.	Click **Close**.	
15.	When the update and reset finish successfully, repeat this procedure, starting with Step 2, for the SP in slot 55 of this enclosure.	

14. Configure Topology of Enclosure If Necessary

See the target system topology (Tetra 8 or Tetra 16) that you recorded on Worksheet 1 under Current Topology. Perform these steps to configure the same topology on the enclosure you are adding:

△ **Caution.** It is critical that you set the topology value correctly on each enclosure you add. This value must always match the value on the target system.

Step	Action	Result
1.	In the initial OSM or TSM Low-Level Link window (before system discovery), click the **System Actions** button.	The Attributes box appears.
2.	Click **Show actions.**	The Actions dialog box appears.
3.	Select the correct Topology to match the target system topology:	
	Select	**Target System Topology**
	Set Topology to Tetra 16	Tetra 16
	Set Topology to Tetra 8	Tetra 8
4.	Click **Perform Action.**	The Status listing tracks the action.
5.	When the action is listed as Complete, click **Close.**	
6.	Click **System Discovery.**	When the Management window appears, the Topology value in the details pane should indicate the change.

Plan for Possible Recabling

When you attach this enclosure to the target system, you might have to recable and reconfigure it. Whether you must do this depends on:

- The topology of your target system
- Where you are adding this enclosure in that topology

An enclosure probably must be recabled or reconfigured if you add it to a Tetra 16 topology as any group in the sequence 05 through 08 or 51 through 89.

This possibility of recabling and reconfiguring means that you must pay careful attention to Worksheets 3 and 4 for the Tetra 8 topology or Worksheets 5 and 6 for the Tetra 16 topology.

15. Power Off Enclosure

You must power off the enclosure before connecting it to the target system.

Step	Action	Result
1.	From the OSM or TSM Low-Level Link, select the **Display** menu.	A menu appears.
2.	Select **Actions.**	A menu appears.
3.	Select **Power Off.**	
4.	Click **Perform action.**	
5.	Click **OK.**	
6.	Verify that the enclosure is powered off.	

16. Repeat Steps 6 Through 15 If Necessary

Skip this step if you are creating a block that contains only one enclosure:

If You Are Creating ...	Then ...
A block containing only one enclosure	Skip this step. Go to Step 18. Change Group Numbers of Block to Fit Target System on page 6-27.
A block containing more than one enclosure	Perform this step.

Configure each enclosure separately before you assemble them into a block. This practice greatly reduces troubleshooting.

If you need to move the system console between enclosures, see Move the System Console on page A-8.

To prepare each enclosure for configuration into a block:

17. Assemble Enclosures Into a Block

Skip this step if you are creating a block that contains only one enclosure:

If You Are Creating …	Then …
A block consisting of only one enclosure	Skip this step and go to Step 18. Change Group Numbers of Block to Fit Target System on page 6-27.
A block consisting of more than one enclosure	Perform this step.

In this step, you attach power-on cables and ServerNet cables to all enclosures. You do not attach any cables between the block and the target system yet.

17.1 Cable Enclosures Together

1. Temporarily set the enclosure numbers in each enclosure of the block according to these rules:

Block Component	Group Number
Processor enclosure	Group 01
I/O enclosures	Groups 11, 12, 13, 14, 15, and so on, in sequential and ascending order, up to the maximum supported by the target system

   To set the enclosure numbers, complete Step 8. Change Group Number of Enclosure to 01 on page 6-16.

2. Connect power-on cables between the enclosures of the block. Power-on cables should be attached so that each PMF CRU and IOMF CRU can send the power-on signal to the next. For cabling patterns, see the *NonStop S-Series Hardware Installation and FastPath Guide*.

3. Connect ServerNet cables for a multiple-enclosure block according to Figure 6-3.

Figure 6-3. ServerNet Cable Connections for a Multiple-Enclosure Block

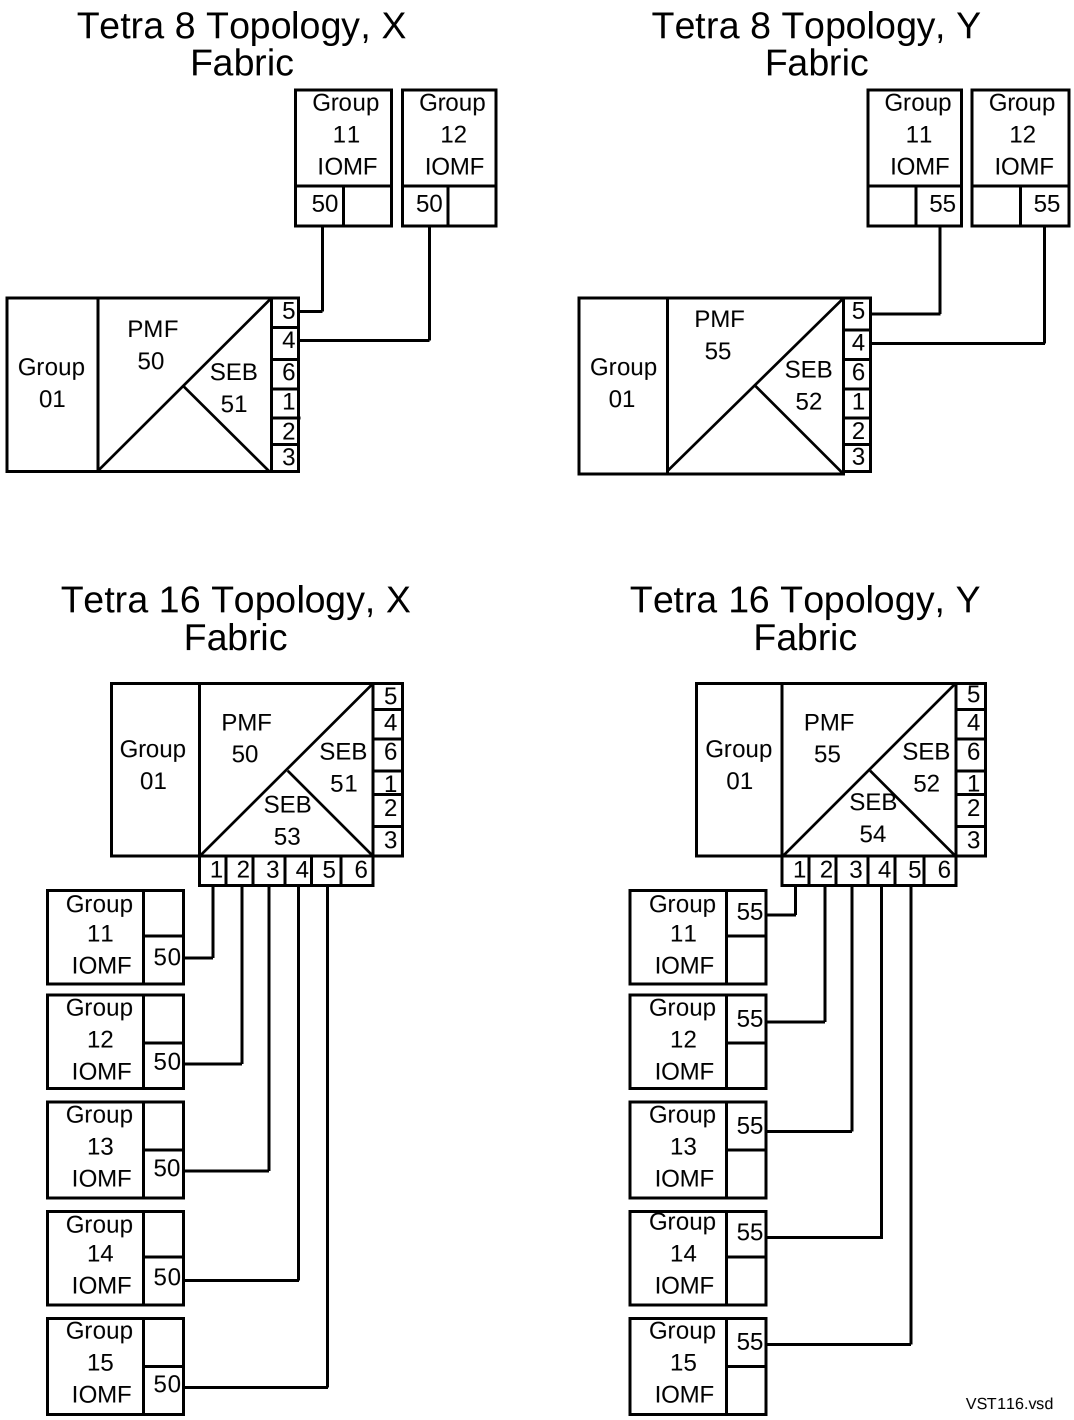

△ **Caution.** ServerNet cabling for IOAM enclosures is different from the cabling shown in this diagram. Only service providers trained by HP should perform cabling of IOAM enclosures.

4. Connect the system console to the block. You performed these procedures for each of the enclosures you have prepared. Now you perform these procedures for the entire block.

17.2 Verify ServerNet Cabling

To be sure that enclosures within the block are cabled correctly, verify the ServerNet cabling by performing this step for each PMF CRU and IOMF CRU in the block:

Step	Action	Result
1.	From the OSM or TSM Low-Level Link, select the Display menu and click **SP Actions.**	The Service Processor Actions dialog box appears.
2.	From the Service Processor (SPs) list, select an SP for an enclosure in the block you are adding.	
3.	From the Actions list, select **SP Responsive Test.**	
4.	Click **Perform action.**	When the action is completed, the status is listed as Completed in the Action Status box.
5.	Repeat this procedure from Step 2 for the other PMF CRU or IOMF CRU in the enclosure.	
6.	Repeat this procedure from Step 2 for every other enclosure in the block you are adding.	

Note. For IOAM enclosures, your service provider must verify the ServerNet cabling by using the ME Responsive Test of the OSM Low-Level Link Application.

17.3 Power Off Block

Power off the block, as described in 15. Power Off Enclosure on page 6-23.

18. Change Group Numbers of Block to Fit Target System

In Step 8. Change Group Number of Enclosure to 01 on page 6-16, you temporarily changed the group numbers of each enclosure in the block to configure them. Now that you are adding the block to the target system, you must change those temporary group numbers to working group numbers appropriate for the target system:

1. See Worksheets 3 and 4 (for a Tetra 8 topology) or 5 and 6 (for a Tetra 16 topology) to see which enclosures receive which group numbers.
2. Change the group numbers by using the group ID switches. For the procedure, see 8. Change Group Number of Enclosure to 01 on page 6-16.

Note. To change the group number of IOAM enclosures, your service provider uses the Configure Module action in the OSM Low-Level Link Application.

---

19. Disconnect System Console From Block

19.1 Stop OSM or TSM Low-Level Link

1. From the OSM or TSM Low-Level Link menu bar, select **File**.
2. Click **Exit**.

19.2 Disconnect System Console

Disconnect the Ethernet cables that connect the system console to the PMF CRUs in slots 50 and 55 of the group 01 enclosure. Figure 6-4 on page 6-28 shows an example of the location of these cables. Some PMF CRUs have the Ethernet port in a different location than the PMF CRU in this example.

Figure 6-4. Disconnecting Ethernet Cable

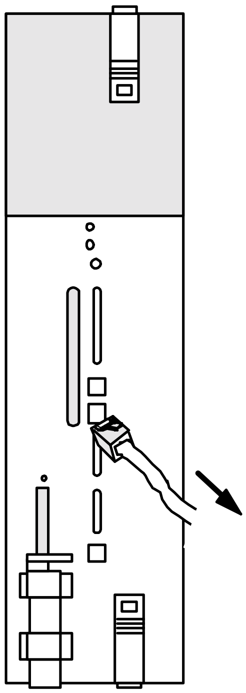

20. Power On Added Block

On each PMF CRU or IOMF CRU in the added block, press and hold the power-on button for two seconds before releasing it. The location of the power-on button is shown in Figure 6-5 on page 6-29.

Pressing the power-on button on one PMF CRU or IOMF CRU causes the other PMF CRU or IOMF CRU in the same enclosure to power on as well. The power-on cables between the enclosures of the added block then carry the power-on signal to the rest of the enclosures of the block.

All CRUs in the block you just added should start automatically. They might take several minutes to run their power-on self-tests (POSTs). Verify that all CRUs are operating:

- The amber LEDs on disk drives and on CRUs installed in slots 50 through 55 become lit briefly as these CRUs are brought up. However, if any of these amber LEDs remain lit for more than five minutes, see POST Fails on page 7-9.
- Green LEDs should be lit.
- The fans on the appearance side of the enclosures of the block should be working.

  If ambient noise prevents you from hearing whether the fans are working check for air flow by holding your hand at the grill located just above slots 51, 52, 53, and 54.

After applying power to the enclosures, if the LEDs on a PMF CRU or IOMF CRU remain unlit, see CRU Does Not Start on page 7-6.

Figure 6-5. PMF/IOMF CRU Power On Button

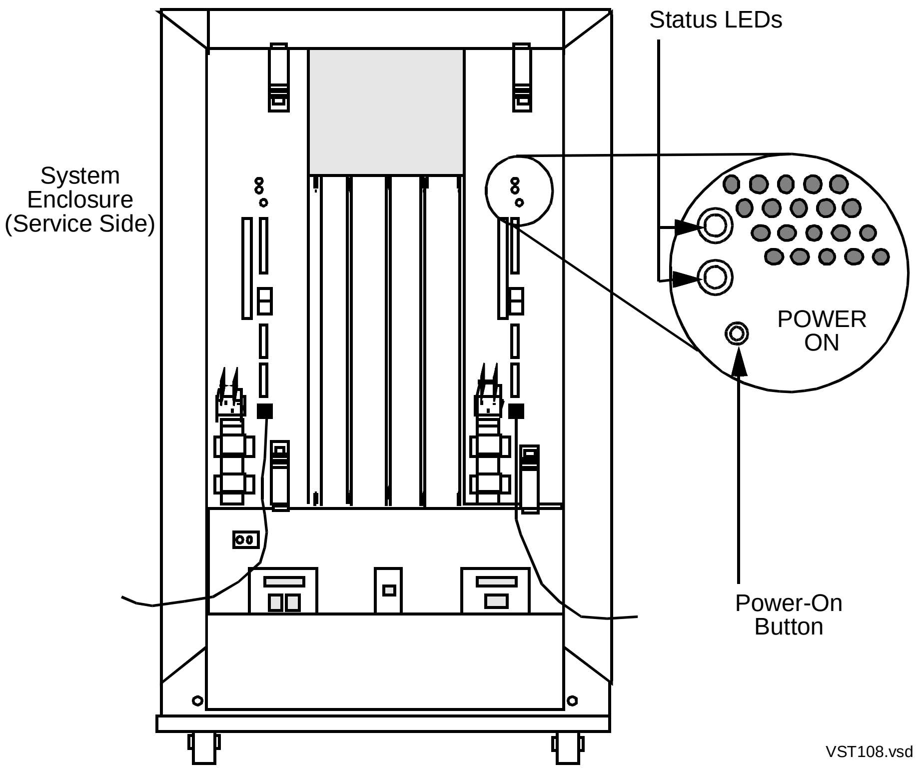

21. Cable Block to Target System

△ **Cautions.** After the AC power cord on the target system or any enclosure in the block is connected to an AC power source, do not unplug and replug it. If you do, see AC Power Cord Was Unplugged or Replugged on page 7-2.

21.1 Install Components on Target System If Necessary

The target system might require components such as extra SEBs or MSEBs to accommodate the block you are adding. Install these components now.

For Information About	See
Which components your system requires	*NonStop S-Series Planning and Configuration Guide*
Installing most system components	CSSI Web on page xiii
Installing SEBs or MSEBs	Guided Replacement Procedures on page xiv
Installing 6760 ServerNet/DAs	Guided Replacement Procedures on page xiv
Installing ServerNet adapters other than the 6760 ServerNet/DA	The manual that discusses that adapter

21.2 Connect ServerNet Cables Between Block and Target System

△ **Caution.** When you attach the block to the target system, you must attach all ServerNet cables correctly the first time.

All ServerNet cabling connections appear in Appendix B, ServerNet Cabling, and on Worksheets 3, 4, 5, and 6 in Appendix C.

You might have to recable the block from the configuration you gave it in Step 17. Assemble Enclosures Into a Block on page 6-24. Whether you must recable depends on:

- The topology of your target system
- Where you are adding this enclosure in that topology

An enclosure probably must be recabled or reconfigured if you add it to a Tetra 16 topology as any group in the sequence 05 through 08 or 51 through 89.

When you attach the ServerNet cables, tighten the jack screws at either end.

21.3 Connect Power-On Cables Between Block and Target System

For Information About	See
Cabling patterns	*NonStop S-Series Hardware Installation and FastPath Guide*

22. Verify Resized Target System

In this step, you run two validation checks on the resized target system.

22.1 Verify That Added Enclosures Appear in System Configuration

If the system console attached to the target system is running the OSM or TSM Low-Level Link:

Step	Action	Result
1.	Log off of the OSM or TSM Low-Level Link.	These actions synchronize the OSM or TSM Low-Level Link with the resized system.
2.	Log back on.	
3.	From the OSM or TSM Low-Level Link window, click **System Discovery.**	The Management window appears. The display might take several minutes to finish.
4.	In the tree pane, located on the left side of the window, double-click the icon for the system.	All enclosures, including the enclosures just added, should be displayed. The display might take several minutes to finish. If all enclosures do not become visible after several minutes, see System Discovery Does Not Show All CRUs or Enclosures on page 7-13.
5.	Log on to the OSM Service Connection or TSM Service Application.	The Management window appears. All enclosures, including the enclosures just added, should be displayed. The display might take several minutes to finish. If all enclosures do not become visible after several minutes, or if red or yellow Abnormal State icons appear, see System Discovery Does Not Show All CRUs or Enclosures on page 7-13.

22.2 Check Configuration of All CRUs

In this step, check whether the CRUs in the target system were configured correctly and whether they passed their power-on self-tests (POSTs).

1. Start a TACL session and log on as a super-group user.
2. Start SCF and enter this command:

```
->STATUS ADAPTER $ZZSTO.*
```

The command returns listings for all storage adapters in the target system. The following example shows what a partial display might look like. In this example, the information indicating that POSTs for newly added enclosure 21 have failed appears in bold type.

```
STORAGE - Status ADAPTER \SIERRA.$ZZSTO.#IOMF.GRP-11.MOD-1.SLOT-50
 Location   Status   POST    Power-1 Power-2 SACs
 (11,1,50)  PRESENT  PASSED  ON      ON      3

STORAGE - Status ADAPTER \SIERRA.$ZZSTO.#IOMF.GRP-11.MOD-1.SLOT-55
 Location   Status   POST    Power-1  Power-2  SACs
 (11,1,55)  PRESENT  PASSED  ON       ON       3

STORAGE - Status ADAPTER \SIERRA.$ZZSTO.#IOMF.GRP-21.MOD-1.SLOT-50
 Location   Status   POST    Power-1  Power-2  SACs
 (21,1,50)  ABSENT   FAILED  OFF      OFF      0

STORAGE - Status ADAPTER \SIERRA.$ZZSTO.#IOMF.GRP-21.MOD-1.SLOT-55
 Location   Status   POST    Power-1  Power-2  SACs
 (21,1,55)  ABSENT   FAILED  OFF      OFF      0
```

All the enclosures in the block you just added should display POST PASSED for slots 50 and 55. If you receive any other message, see POST Fails on page 7-9.

23a. Update Firmware and Code in Block (Using TSM)

Do this step even if you performed 13. Update SP Firmware in Enclosure If Necessary.

For each PMF CRU or IOMF CRU in each enclosure you added, check (and update if necessary) the SP firmware, processor boot code, and SCSI boot code:

Step	Action	Result
1.	In the TSM Service Application, select **Display.**	A menu appears.
2.	Select **Single SP Firmware Update.**	The Single SP Firmware Update dialog box appears.
3.	From the Service Processor CRU listing, select the PMF CRU in slot 50 of the processor enclosure you just added.	These selections are displayed: ● Service processor firmware ● Processor boot code ● SCSI boot code
4.	Examine the State Box for the Processor boot code. ● If the State Box is OK, skip to Step 5. ● If the State Box is not OK, perform Step 4a through Step 4d:	
a.	Click the check box for the Processor boot code.	
b.	Click **Perform action.**	The Update Action Summary dialog box appears.
c.	When the Status line lists the action as Completed, click **Close.**	The Update Action Summary dialog box disappears.
d.	Perform a hard reset of the processor: (Do not close the TSM Service Application session; you will use it again in Step 5.)	
	1. Launch and log on to the TSM Low-Level Link.	
	2. Click **Processor Status**.	The Processor Status dialog box appears.
	3. Click to select the processor you want to update.	
	4. From the Processor Actions list, select **Hard Reset** and then click **Perform action**.	After one to three minutes, the Action Status box indicates a successful Hard Reset.
	5. When the status line lists the action as Completed, click **Close**.	If the action does not complete successfully, see Action Detail for more information.
5.	Examine the State Box for the SCSI boot code. ● If the State Box is OK, skip to Step 6. ● If the State Box is not OK, perform Step 5a through Step 5c:	
a.	Click the check box for the SCSI boot code.	

Step	Action	Result
b.	Click **Perform action.**	The Update Action Summary dialog box appears.
c.	When the Status line lists the action as Completed, click **Close.**	The Update Action Summary dialog box disappears.
6.	Examine the State Box for the SP Firmware. ● Skip to Step 15 if both of the following conditions are true: ● The state box is OK ● You performed 13. Update SP Firmware in Enclosure If Necessary. ● Otherwise, perform Step 6a through Step 6c:	
a.	Click the check box for the SP Firmware.	
b.	Click **Perform action.**	The Update Action Summary dialog box appears.
c.	When the Status line lists the action as Completed, click **Close.**	The Update Action Summary dialog box disappears.
7.	Click **Cancel.**	The Single SP Firmware Update dialog box disappears.
8.	In the Physical view, double-click the processor enclosure for the block you just added.	The Physical view displays components on both sides of the enclosure.
9.	Right-click the slot for the SP you just updated.	A menu appears.
10.	Select **Actions**.	The Actions dialog box appears.
11.	From the Select Component box, select **Service Processor.**	Choices appear in the Actions for list.
12.	In the Actions for list, select **Reset.**	
13.	Click **Perform action.**	After one to three minutes, the Action Status box indicates a successful Reset.
14.	When the status line lists the action as Completed, click **Close.** If the status line never lists the action as Completed, see the TSM documentation.	
15.	For the PMF CRU in slot 55 of the enclosure you added, repeat this procedure from Step 3.	
16.	If you added any other enclosures, repeat this procedure for each enclosure from Step 1.	

23b. Update Firmware and Code in Block (Using OSM)

Do this step even if you performed 13. Update SP Firmware in Enclosure If Necessary.

Note. For IOAM enclosures, your service provider must verify that the ME firmware and the ME FPGA is the latest version. Updating the ME firmware and ME FPGA is described in the *Modular I/O Installation and Configuration Guide*.

For each PMF CRU or IOMF CRU in each enclosure you added, check (and update if necessary) the SP firmware, processor boot code, and SCSI boot code:

Step	Action	Result
1.	In the OSM Service Connection, from the **Display** menu, select **Multi-Resource Actions**.	The Multi-Resource Actions dialog box appears.
2.	From the **Resource Type** drop-down list, select **SP**.	
3.	Select **Processor Boot Millicode Firmware Update** from the **Action** drop-down list.	
4.	In the list of SPs, scroll to find the Processor Boot Millicode Firmware attributes and compare the **Version** attribute to the **Default File Version**. ● If the **Version** is older than the **Default File Version**, the **Version** is down-rev and should be updated. Perform Steps 4a through 4e. ● If the **Version** is the same as or newer than the **Default File Version**, skip to Step 5.	
a.	Select the SP you need to update and click **Add**.	The SP moves to the lower list.
b.	Click **Perform Action**.	After a confirmation dialog, a Parameter Input dialog box allows you to specify an alternate file location for the version to use for the update.
c.	Enter a location or leave it blank if you are using the default location, and click **OK** to perform the update.	The Progress bar at the bottom of the Multi-Resource Actions dialog box indicates the pass or fail status of the action.
d.	When the action completes, check the Processor Boot Millicode Firmware attributes again to confirm that the **Version** now matches the **Default File Version**. If the action fails, click **Action Summary** and then **Details** to get more information about why the action failed.	
e.	Perform a hard reset of the processor (do not close the OSM Service Connection session or the Multi-Resource Actions dialog box; you will use it again in Step 5):	

Step	Action	Result
	1. Launch and log on to the OSM Low-Level Link.	
	2. Click **Processor Status**.	The Processor Status dialog box appears.
	3. Click to select the processor you want to update.	
	4. From the Processor Actions list, select **Hard Reset** and then click **Perform action**.	After one to three minutes, the Action Status box indicates a successful Hard Reset.
	5. When the status line lists the action as Completed, click **Close**. If the action does not complete successfully, see Action Detail for more information.	
5.	Return to the Multi-Resource Actions dialog box. Select **SCSI Boot Code Firmware Update** from the **Action** drop-down list.	
6.	Check the SCSI Boot Code Firmware attributes for the same SP and compare the **Version** attribute to the **Default File Version**. ● If the **Version** is older than the **Default File Version**, the **Version** is down-rev and should be updated. Perform Steps 6a through 6e. ● If the **Version** is the same as or newer than the **Default File Version**, skip to Step 7.	
a.	If the SP you want to update is not already in the lower list, select the SP and click **Add** to move it to the lower list.	The SP moves to the lower list.
b.	Click **Perform Action**.	After a confirmation dialog, a Parameter Input dialog box allows you to specify an alternate file location for the version to use for the update.
c.	Enter a location or leave it blank if you are using the default location, and click **OK** to perform the update.	The Progress bar at the bottom of the Multi-Resource Actions dialog box indicates the pass or fail status of the action.
d.	When the action completes, check the Processor Boot Millicode Firmware attributes again to confirm that the **Version** now matches the **Default File Version**. If the action fails, click **Action Summary** and then **Details** to get more information about why the action failed.	
7.	For the PMF CRU in slot 55 of the enclosure you added, repeat this procedure from Step 3.	
8.	If you added any other enclosures, repeat this procedure for each enclosure from Step 1.	

24. Reload Processors in Block If Necessary

Skip this step if the added block does not contain a processor enclosure.

If the Added Block Contains …	Then …
One I/O enclosure and no processor enclosures	Skip this step. Go to 25. Verify Operations in Added Block on page 6-38.
A processor enclosure, with or without I/O enclosures	Perform this step.

Note. Ensure you have KMSF swap files of the proper size and location available on the donor system after the reduction is finished. Otherwise, applications and processes on the newly resized system might terminate. For information about KMSF swap files and their recommended sizes, see the *Kernel-Managed Swap Facility (KMSF) Manual*.

1. Start a TACL session and log on as a super-group user.
2. At the TACL prompt:

   `>RELOAD (`*`m,n`*`), PRIME`

   where *m* and *n* are the processor numbers in the block you just added.

   Processor numbers are assigned according to group number:

In group …	The processors are …	In group …	The processors are …
01	0,1	05	8,9
02	2,3	06	10,11
03	4,5	07	12,13
04	6,7	08	14,15

   In this example, you reload the processors in group 04:

   ```
   TACL 1>RELOAD (6,7), PRIME
   ```

25. Verify Operations in Added Block

This step ensures that all enclosures in the added block are operating normally:

Step	Action	Result
1.	Log on to the OSM Service Connection or TSM Service Application.	The Management window appears.
2.	In the tree pane, click the appropriate icons to view each enclosure in the added block.	All added enclosures and their components should be displayed. The display might take several minutes to finish. If all added enclosures with their components do not become visible after several minutes, see System Discovery Does Not Show All CRUs or Enclosures on page 7-13.
3.	If any warning triangles or other signs of abnormality are displayed, diagnose and solve the problems. See the OSM or TSM online help.	

26. Configure CRUs in Added Block

Now that you have connected the added block to the target system, you must configure the CRUs in each enclosure of that block.

Finish configuring all the CRUs in one enclosure before moving to the next enclosure in the block. Configuring CRUs one enclosure at a time minimizes variables in troubleshooting.

How you configure your CRUs depends on your system and subsystems. System-specific and subsystem-specific configuration information appears in the manuals devoted to those systems and subsystems.

7 Troubleshooting

This section gives troubleshooting information to solve problems encountered during system expansion or reduction.

Note. For problems with components in an IOAM enclosure, your service provider should see the *Modular I/O Installation and Configuration Guide*.

Symptoms and Problems

AC Power Cord Was Unplugged or Replugged

Cause

You unplugged the enclosure's AC power cords from the power source and then plugged them back in. This might have occurred during Step 9. Power On Enclosure on page 6-17 or Step 20. Power On Added Block on page 6-28.

Recovery

1. Unplug the AC power cords of the affected block or enclosure.
2. Unplug the battery packs from the power monitor and control unit (PMCU) in the affected block or enclosure.
3. Replug battery packs into the PMCU in the affected block or enclosure, but do not plug in the AC power cords yet.
4. Begin the resizing process again at the step where you left off:

Cause	Return to ...
AC Power Cord Was Unplugged or Replugged	Depending on what step brought you here, either Step 9. Power On Enclosure on page 6-17 or Step 20. Power On Added Block on page 6-28
Added Enclosure Does Not Power On	Step 9. Power On Enclosure on page 6-17
CRU Does Not Start	Step 9. Power On Enclosure on page 6-17

Added Enclosure Does Not Power On

Not all enclosures in the added block powered on after you pressed the power-on button, possibly from Step 9. Power On Enclosure on page 6-17.

Cause and Recovery

Cause	Recovery
AC power cords are unplugged.	See AC Power Cord Was Unplugged or Replugged on page 7-2.
Power-on cables are disconnected.	Figure 7-1 shows the location of the power-on cables. The *NonStop S-Series Hardware Installation and FastPath Guide* describes the procedure for adding power-on cables between enclosures. When you add power-on cables to enclosures, each enclosure must have a way to send the power-on signal to the next enclosure. After you verify the power-on cabling, press power-on buttons again in all nonfunctional PMF CRUs and IOMF CRUs. Return to Step 9. Power On Enclosure on page 6-17.

Figure 7-1. PMF/IOMF CRU Power-On Cables

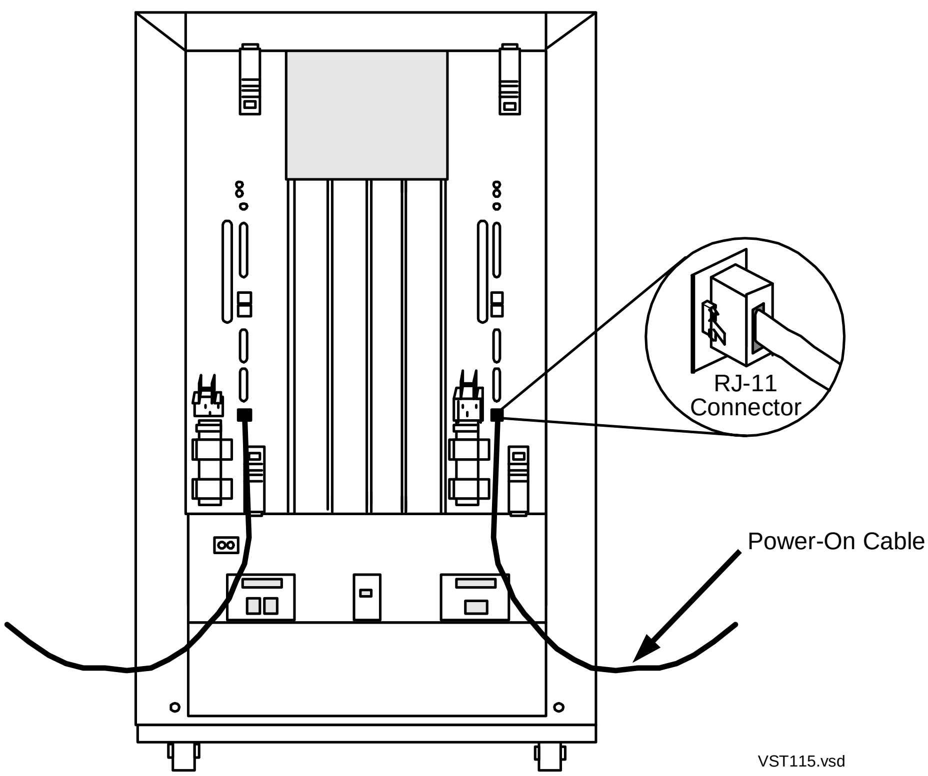

Added Enclosure Is Not Recognized by the Target System

The OSM or TSM applications on the target system do not show the enclosure you just added. This problem might occur during Step 22.1 Verify That Added Enclosures Appear in System Configuration on page 6-31.

Cause and Recovery

Cause	Recovery
Topology problem. The topology value for the added block, as specified in the OSM or TSM Low-Level Link, does not agree with your target system's value.	1. Unplug the AC power cords of the affected block or enclosure. 2. Unplug the battery packs from the power monitor and control unit (PMCU) in the affected block or enclosure. 3. Replug the battery packs into the PMCU in the affected block or enclosure, but do not plug in the AC power cords yet. 4. Attach a system console. See Step 7 on page 6-15. 5. Change the group ID to 01. See Step 8 on page 6-16. 6. Power on the enclosure. See Step 9 on page 6-17. 7. Correct the topology value of the enclosure. See Step 14 on page 6-22. 8. Resume the expansion procedure at Step 15. Power Off Enclosure on page 6-23.
ServerNet fabrics are cross-cabled.	See ServerNet Fabrics Are Cross-Cabled on page 7-10. Resume the expansion procedure at Step 20. Power On Added Block on page 6-28.
AC power cords are unplugged.	See AC Power Cord Was Unplugged or Replugged on page 7-2.
Power-on cables are disconnected.	See Added Enclosure Does Not Power On on page 7-3.
The enclosure has an incorrect group ID switch.	1. Unplug the AC power cords of the affected block or enclosure. 2. Unplug the battery packs from the power monitor and control unit (PMCU) in the affected block or enclosure. 3. Replug battery packs into the PMCU in the affected block or enclosure, but do not plug in the AC power cords yet. 4. Resume the expansion procedure at Step 18. Change Group Numbers of Block to Fit Target System on page 6-27.
Malfunctioning hardware.	Contact your service provider.

If the preceding recovery procedures do not work, contact your service provider.

Amber LEDs Remain Lit

Green LEDs should remain lit on powered-on enclosures. However, if an amber LED remains lit, see POST Fails on page 7-9.

Neither CRU Starts

See Added Enclosure Is Not Recognized by the Target System on page 7-4.

Cannot See Enclosure in OSM or TSM

See Added Enclosure Is Not Recognized by the Target System on page 7-4.

Configuration File Already Exists

When you are updating the system configuration database during one of these steps:

- 9. Finish the Reduction on page 4-19
- 1. Prepare Target System for Expansion on page 6-4
- 3. Prepare Target System for Addition of Block on page 6-9

you get this message:

```
Configuration file $SYSTEM.ZSYSCONF.CONFxxyy already exists.
Replace it? (Y/[N])
```

Cause

You requested that SCF save the configuration file with a name that already exists. If you answer Y, you overwrite that existing system configuration database.

Recovery

Enter N and retry the SAVE CONFIGURATION *xx.yy* command with another value for *xx.yy*.

CRU Does Not Start

In Step 9. Power On Enclosure on page 6-17, when you press the power-on button on one PMF CRU or IOMF CRU, the other PMF CRU or IOMF CRU in the same enclosure does not power on.

Cause

The PMF CRUs or IOMF CRUs in a single enclosure might not be seated correctly.

Recovery

1. Unseat the affected CRU and wait one minute.
2. Reseat the CRU and reattach the AC power cord within 25 seconds.

   The CRU performs a POST. If an amber LED stays on for longer than five minutes, see POST Fails on page 7-9.

If the problem persists, contact your service provider.

IP Address Already Exists

During Step 11. Configure System Console and Enclosure on page 6-19, when you attempt to add the IP addresses of the enclosures in the block in the OSM or TSM Low-Level Link, you get this message:

```
You have entered IP addresses that already appear in the current
system list. Do you want to continue?
```

This is a normal message.

Cause

The enclosures in the block to be added have the same low-level link IP address (as with new enclosures, for example).

Recovery

1. In answer to the question, click **Yes.**
2. The Edit Systems List dialog box then displays the added enclosure with the name “New System.” Click **OK**.
3. In the OSM or TSM Low-Level Link, log on to New System. The new system entry is no longer listed in the Edit Systems List dialog box.

   If the system console you are using is the system console and you are moving it from the target system, the donor system, or the enclosure, you might have to re-add the IP address later when you reattach the system console to the target system.

Ping Command Does Not Work

The ping command on the system console (entered during Step 10. Verify Connection Between System Console and Enclosure on page 6-18) returns the message "Destination host unreachable" or "Request timed out."

Cause

You cannot connect to the block or enclosure you are configuring because:

- There is a disconnect somewhere in the Ethernet cables.
- The subnet mask on the system console does not match the subnet mask on the enclosure or block.
- Your terminal might have an IP address that is not compatible with the network.

Recovery

You must be able to communicate with the SPs in the enclosure. Therefore, your IP address must be in the same subnet as the SPs. If the enclosure is using the default IP addresses, your IP address must be 192.231.36.*n*, where *n* is some unused number between 0 and 225.

1. Verify that all cables and Ethernet switches (or hubs) between the system console and the block or enclosure are properly connected and fully functional.
2. Verify the IP configuration of the system console:

   Display the IP address of the system console:

   a. Select **Start > Programs > Command Prompt**.

   b. Enter this command:

      `ipconfig`

   The response to this command resembles this display:

```
C:\users\default> ipconfig

Windows NT IP Configuration

Ethernet adapter NdisWan4:

      IP address. . . .. . . . . : 155.186.250.125
      Subnet Mask .  . . . . . . : 255.255.0.0
      Default Gateway .  . . . . : 155.186.250.124

C:\users\default>
```

3. Compare the Subnet Mask and Default Gateway values in the display with the IP configuration of the enclosure that you recorded on Worksheet 7.

4. If these values differ, change the IP address of the system console:

   a. Click **Start > Settings > Control Panel.**
   b. Double-click the **Network** icon.
   c. Click the **Protocols** tab.
   d. Select **TCP/IP Protocol**.
   e. Click **Properties**.
   f. In the dialog box, click the **IP Address** tab.
   g. If necessary, click **Specify an IP address**.
   h. Change the IP Address, Subnet Mask, and Default Gateway fields:

      - IP address: specify a valid IP address not currently used in the network. Factory-assigned default IP addresses appear in the following figure.
      - Subnet Mask: the Subnet Mask entry from Worksheet 7.
      - Default Gateway: the Gateway Route entry from Worksheet 7.

   i. Click **OK.**
   j. Click **Close**.

5. Restart the system console.

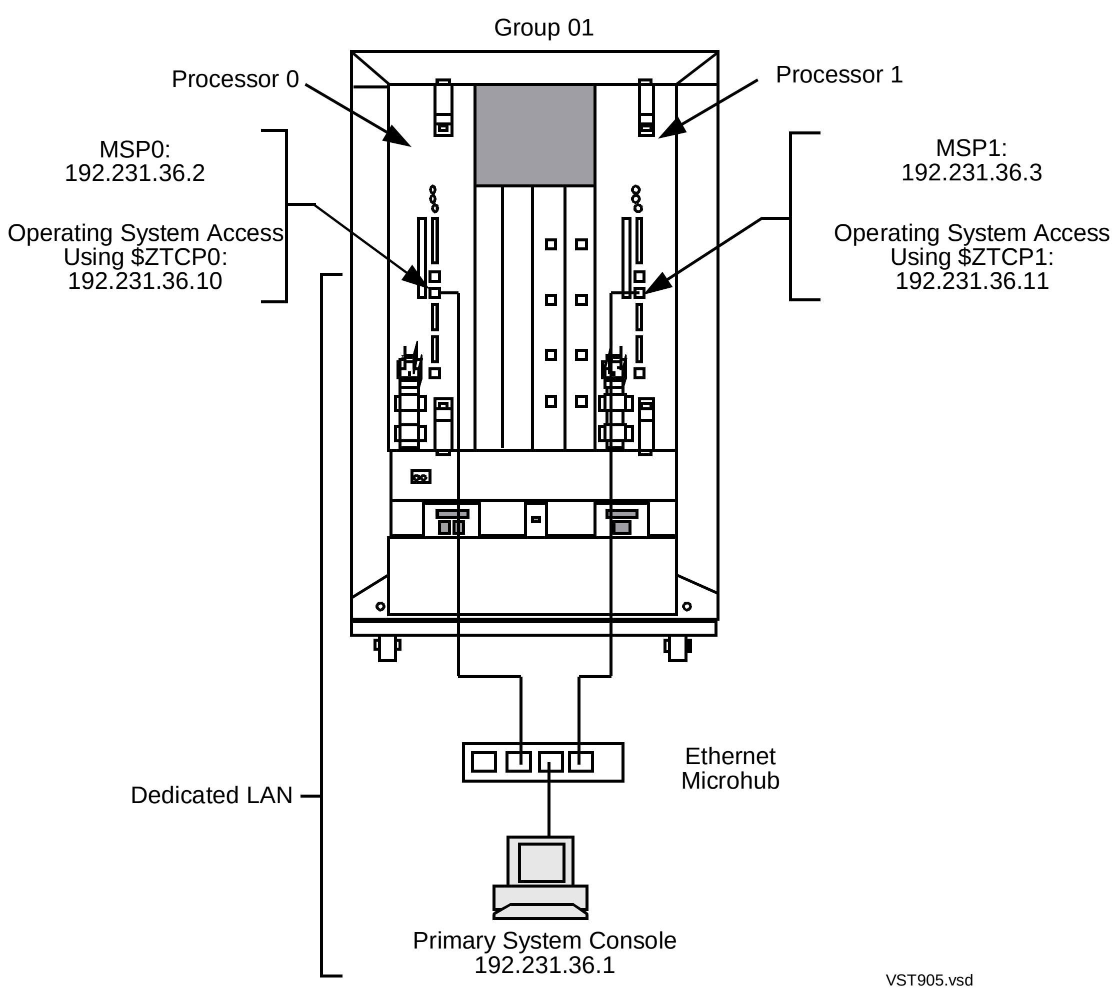

POST Fails

The power-on self-test (POST) fails for one or more PMF CRUs or IOMF CRUs. This situation might occur when you perform these steps:

- 9. Power On Enclosure on page 6-17
- 20. Power On Added Block on page 6-28
- 22.2 Check Configuration of All CRUs on page 6-32.

Note. When POST fails before the added block is connected to the target system, the LEDs are available but not the event messages.

Symptoms of POST failure include:

- The amber LED on the CRU remains lit for more than five minutes after the POST finishes
- SP event 602 (ZSPR-EVT-CRU-POST-FAILED)
- Storage event 1033 (ZSTO-EVT-CRU-INIT-FAILED) in the $ZLOG service log
- Storage event 1035 (POST of the CRU failed)

Symptoms of a successful POST include:

- SP event 601 (ZSPR-EVT-CRU-POST-SUCCEEDED)
- Storage event 1032 (ZSTO-EVT-CRU-INIT)
- Storage event 1034 (POST of the CRU succeeded)

Cause

A CRU has failed.

Recovery

Contact your service provider.

RELOAD Command Does Not Work

Not all processors in the system come up in response to a TACL RELOAD command, as might happen during Step Skip this step if the added block does not contain a processor enclosure. on page 6-37.

Cause

The ServerNet X fabric might be down.

Recovery

1. Enter this command and check for any status other than UP.

```
> SCF STATUS SERVERNET $ZSNET
```

2. Start any down fabrics:

```
> SCF START SERVERNET $ZSNET.X.*
> SCF START SERVERNET $ZSNET.Y.*
```

3. Return to Step Skip this step if the added block does not contain a processor enclosure. on page 6-37.

If the problem persists, contact your service provider.

Request Timed Out

See Ping Command Does Not Work on page 7-7.

ServerNet Fabrics Are Cross-Cabled

Among the symptoms of this problem is an inability to see the enclosure in the OSM or TSM Low-Level Link.

Cause

You attached one or more X fabric ServerNet cables to the Y fabric or vice versa. This miscabling might have occurred during either of the following steps:

- Connecting the cables wrong in 21.2 Connect ServerNet Cables Between Block and Target System on page 6-30
- Labeling the cables wrong in Step 5 on page 6-5

Recovery

1. Disconnect ServerNet cables between the target system and the affected block.
2. Unplug the AC power cords of the affected block.
3. Unplug the battery packs in the affected block.
4. Replug battery packs in the affected block.
5. Begin the resizing process again at 20. Power On Added Block on page 6-28.

ServerNet Fabrics Are Down

One or more ServerNet fabrics might be down.

Cause

This can happen either:

- On a system running a RVU earlier than G06.04
- If a fabric was brought down manually

Recovery

1. Enter the following command and check for any status other than UP.

```
> SCF STATUS SERVERNET $ZSNET
```

2. Start any down fabrics:

```
> SCF START SERVERNET $ZSNET.X.*
> SCF START SERVERNET $ZSNET.Y.*
```

3. Return to Step 3.2 Eliminate Problems in the ServerNet Fabrics on page 4-6.

If the problem persists, contact your service provider.

Slots 50 and 55 Remain Unlisted for an Enclosure

After you connect the block to the target system, an SCF STATUS ADAPTER command does not list slots 50 and 55 for one or more enclosures in the block.

Cause and Recovery

See Added Enclosure Is Not Recognized by the Target System on page 7-4.

Code 316 Returned During Logon

While trying to log on to the OSM or TSM Low-Level Link during Step 12.1 Determine SP Firmware Version for Enclosure on page 6-20, you get a code 316 error message on the status bar at the bottom of the OSM or TSM window.

Cause

You are trying to configure an enclosure or block before adding it to a target system, but the system console does not recognize the enclosure you are trying to configure because you are connected to a group other than group 01.

Recovery

1. Unplug the AC power cords of the affected block or enclosure.
2. Unplug the battery packs from the power monitor and control unit (PMCU) in the affected block or enclosure.
3. Replug battery packs into the PMCU in the affected block or enclosure, but do not plug in the AC power cords yet.
4. Return to Step 8. Change Group Number of Enclosure to 01 on page 6-16. Perform this step and all subsequent steps again.

Startup TACL Window Does Not Open

You cannot open a startup TACL window using the procedure Start a Startup TACL Session on page A-9.

Cause and Recovery

	Cause	Recovery
1.	You get a “duplicate CLCI request” message, indicating that a startup TACL window is already open.	This message also gives you the IP address of the workstation that currently owns the startup TACL, from which you can find the workstation.
2.	The startup TACL window has no TACL prompt because you attempted to open a startup TACL window on the wrong ServerNet fabric.	This is the backup fabric window. Minimize it to see the primary fabric window.

If the preceding recovery procedures do not work, contact your service provider.

Storage Automatic Configuration Does Not Work

The AUTOCONFIG option of the SCF ALTER command does not automatically configure all internal disk drives in an enclosure.

An AUTOCONFIG command failure has three possible causes, each with its own recovery procedure.

Note. If AUTOCONFIG fails for any reason, you must configure all disk drives in the enclosure manually with the SCF ADD DISK command. The recovery actions and procedures given here do not eliminate the need to perform this manual configuration. Instead, these procedures are intended to prevent the problem from happening again in the future.

Cause and Recovery

Cause	Recovery
1. A failure occurred during disk configuration	For more information about the failure, see the EMS event messages generated by your system. Configure all internal disk drives in the enclosure manually with the SCF ADD DISK command.
2. The storage subsystem manager process has the Storage Automatic Configuration feature (AUTOCONFIG) set to OFF.	1. Verify this state and set AUTOCONFIG to ON by performing Step 3.2 Check Storage Automatic Configuration and Autostart on page 6-9. This setting takes effect during the next attempt at configuration. 2. Configure all internal disk drives in this enclosure manually with the SCF ADD DISK command.

System Discovery Does Not Show All CRUs or Enclosures

When you click System Discovery in the OSM or TSM Low-Level Link, a CRU (in the block) or enclosure (in the system) is not visible.

See Added Enclosure Is Not Recognized by the Target System on page 7-4.

Topology Does Not Appear

The OSM or TSM Low-Level Link is unable to report the topology value. You might get this result during Step 2.1 Record Basic Information on page 4-4 or Step 2.2 Record Information About the Master Service Processor on page 6-6.

Cause

The topology value continues to display as Retrieving Data.

Recovery

While still in the OSM or TSM Low-Level Link, from the Display menu, select **Refresh**.

Unable to Log On to OSM or TSM Low-Level Link

See Code 316 Returned During Logon on page 7-12.

Service Processor Subsystem Event Messages

The symbolic name of each of these events begins with ZSPR-EVT:

Event Number	Name (ZSPR-EVT)	Description
110	-CRU-INSERTION-OK	CRU insertion passed preliminary self tests.
111	-CRU-INSERTION-FAILED	CRU insertion failed preliminary self tests.
112	-CRU-REMOVED	CRU was removed.
501	-CRU-SNETCFG-STARTED	ServerNet configuration started for unit.
502	-CRU-SNETCFG-FAILED	ServerNet configuration failed for unit.
503	-CRU-SNETCFG-COMPLETE	ServerNet configuration succeeded for unit.
504	-LINK-ALIVE	ServerNet link from router port is now alive.
601	-CRU-POST-SUCCEEDED	Power-on self-test of CRU succeeded.
602	-CRU-POST-FAILED	Power-on self-test of CRU failed.

Storage Subsystem Event Messages

The symbolic name of each of these events begins with ZSTO-EVT:

Event Number	Name (ZSTO-EVT)	**Description** (page 1 of 2)
1026	-ADAPTER-AUTO-FAILED	Autoconfiguration of an adapter failed. This event is usually preceded by storage subsystem event 1025 (failed to access $ZCNF).
1027	-ADAPTER-AUTO-RETRY	Waiting for a resource that is blocking auto-configuration of an adapter.
1028	-CRU-INSERTED	Hardware insertion event sent from the configuration utility process ($ZCNF).
1029	-CRU-REMOVED	Hardware deletion event sent from the configuration utility process ($ZCNF).
1030	-WRONG-CRU-INSERTED	Hardware insertion event received does not match the configuration information in the system configuration database.
1031	-CRU-CONFIGURED	CRU configuration successfully added to the system configuration database.
1032	-CRU-INIT	CRU initialization succeeded.
1033	-CRU-INIT-FAILED	CRU initialization failed. The event contains the error number returned from the support function used. This function is offered by the SCSI Interface Module (SIFM) or by Extensible I/O (XIO).
1034	-POST-COMPLETE	Power-on self-test of CRU succeeded.
1035	-POST-FAILED	Power-on self-test of CRU failed.
1037	-INVALID-CRU	Not enough information available to determine: • What type of CRU was inserted • How to configure that CRU
1040	-UNSUPPORTED-CRU	CRU inserted that is not supported by the storage subsystem.
1041	-CRU-OWNERSHIP	CRU cannot be configured by the storage subsystem. The slot where the CRU was inserted is already configured by another subsystem.
1042	-CRU-DELETED	CRU configuration deleted from the system configuration database.

Event Number	Name (ZSTO-EVT)	**Description** (page 2 of 2)
1045	-CRU-CONFIG-FAILED	Autoconfiguration of CRU failed. This event contains information describing the reason for failure.
1046	-CRU-ACCESS-ERROR	Storage subsystem manager process did not obtain information for CRU.
1047	-CRU-RECONFIG	CRU configuration has been changed.
1048	-CONFIG-DEFERRED	CRU configuration deferred. This event is generated when the storage profile record forces the storage subsystem manager to start an I/O process to complete the configuration of the CRU, but the processors are unavailable where the I/O process is configured to run.
2001	-CNFG-DISK-NAME	Disk has changed name automatically.
2004	-CNFG-CRU-NOT-CONFIG	CRU is not configured. This event is generated only when Storage Automatic Configuration is disabled. In that case, the event is generated when: ● The system is loaded. ● A hardware insertion event is received.
2005	-CNFG-MISSING-CRU	A configured CRU is not present in the system.

A Common System Operations

This appendix describes how to perform common operations involved in resizing your system.

Determine the Processor Type

You can use either the OSM or TSM Low-Level Link or the OSM Service Connection or TSM Service Application to determine the processor type of the NonStop S-series server that you are logged on to. The software does not display the server model number, but you can use the Part Number or Processor Name attributes to infer the server model number.

Determine the Processor Type Using the OSM or TSM Low-Level Link

Step	Action	Result
1.	• In OSM, select **Start>Programs>HP OSM>OSM Low-Level Link**. • In TSM, select **Start>Programs>Compaq TSM>TSM Low-Level Link Application**.	
2.	Log on to a low-level link connection.	
3.	Click **System Discovery** to discover the system.	
4.	From the **Display** menu, select **Physical View**.	
5.	Resize the Management window so that you can view the group numbers under the enclosure icons in the Physical view.	
6.	In the Physical view, double-click the group containing the processor for which you want to determine the processor type.	
7.	Click the PMF CRU containing the processor.	
8.	In the details pane, click the **Attributes** tab.	The attributes appear in the details pane.
9.	Find the Part Number attribute and note its value for the selected PMF CRU.	To determine the processor type and system model number based on this part number, see the OSM or TSM online help or the *NonStop S-Series Planning and Configuration Guide*.

Determine the Processor Type Using the OSM Service Connection or TSM Service Application

Step	Action	Result
1.	• In OSM, select **Start> Programs>HP OSM>OSM Service Connection**. • In TSM, select **Start> Programs>Compaq TSM>TSM Service Application**.	
2.	Log on to the OSM Service Connection or TSM Service Application.	
3.	From the tree pane: a. Double-click the group number for the enclosure. b. Double-click one of the PMF CRUs. c. Select one of the processors.	The attributes for that processor are displayed in the Attributes tab.
4.	From the Details pane, select the **Attributes** tab.	
5.	From the Attributes tab, check the processor name to determine the processor type.	
6.	To determine the system model number for this processor type, see the OSM or TSM online help or the *NonStop S-Series Planning and Configuration Guide*.	
7.	For the other processor in the enclosure, repeat Steps 3b through 6.	

Determine the ServerNet Fabric Status

You use the OSM Service Connection, TSM Service Application, or SCF to determine the status of ServerNet fabrics.

Check the ServerNet Fabric Status Using the OSM Service Connection or TSM Service Application

This procedure is supported by OSM and TSM version 2000A or later and can be used on both processor enclosures and I/O enclosures.

Step	Action	Result
1.	For TSM, in the tree pane, click the system tab (for viewing internal fabrics) or the **Cluster** tab (for viewing external fabrics). For OSM, click the system name.	If the system tab was selected, expand the system icon to find the Internal_SvNet_Fabric_X and Internal_SvNet_Fabric_Y icons. If the Cluster tab was selected, expand first the Cluster icon and then the ServerNet Cluster icon to find the External ServerNet Fabric X and External ServerNet Fabric Y icons.
2.	Right-click the icon for the fabric that you want to check first.	A menu appears.
3.	Select **Actions**.	The Actions dialog box appears.
4.	Click **Perform Action**.	A confirmation dialog box appears.
5.	Select **Group Connectivity ServerNet Path Test**.	
6.	Click **OK**.	It might take several minutes for the test to finish. When it does, the status goes from Started to Completed. If the test fails, proceed to Step 6.
7.	Close the Actions dialog box.	Consult the Repair Actions. See “Viewing Alarm Detail Information” in the OSM Service Connection or TSM Service Application online help. If necessary, contact your service provider
8.	Repeat this procedure for the fabric you want to check next.	

Note. For OSM, windows can be hidden behind the primary window and, if waiting for a response, can create problems. To check for open windows, click the **Window** menu in the OSM toolbar.

Check ServerNet Fabric Status Using SCF (Processor Enclosures Only)

The following example shows how to use the SCF STATUS SERVERNET $ZSNET command to check the ServerNet communications among processors and provide an example of the display that might result. This command shows whether two processors can communicate over a specific fabric.

If the display shows any down paths, see ServerNet Fabrics Are Down on page 7-11.

This SCF command generates listings for the maximum number of processors that are possible in your system, whether or not the system actually contains the maximum number of processors. Ensure that the status for all processors actually in your system is listed as UP.

In this example, processors 0 and 1 cannot communicate with processors 2 and 3 over the Y fabric. Their listings, which indicate a problem, are highlighted in bold type.

Note. Run this ServerNet path test only if the state of the fabric is in question. The OSM and TSM package is designed to keep an up-to-date view of the state of the fabric.

```
$SYSTEM STARTUP 1> scf
SCF - T9082G02 - (16OCT98) (25SEP98) - 02/22/1999 13:37:32 System \ALPHA5
Copyright Tandem Computers Incorporated 1986 - 1998
(Invoking \ALPHA5.$SYSTEM.STARTUP.SCFCSTM)
1-> status servernet $zsnet
NONSTOP KERNEL -  Status SERVERNET
X-FABRIC
   TO     0   1   2   3   4   5   6   7   8   9  10  11  12  13  14  15
FROM
  00      UP  UP  UP  UP  UNA UNA UNA UNA UNA UNA UNA UNA UNA UNA UNA UNA
  01      UP  UP  UP  UP  UNA UNA UNA UNA UNA UNA UNA UNA UNA UNA UNA UNA
  02      UP  UP  UP  UP  UNA UNA UNA UNA UNA UNA UNA UNA UNA UNA UNA UNA
  03      UP  UP  UP  UP  UNA UNA UNA UNA UNA UNA UNA UNA UNA UNA UNA UNA
  04   <- DOWN
  05   <- DOWN
  06   <- DOWN
  07   <- DOWN
  08   <- DOWN
  09   <- DOWN
  10   <- DOWN
  11   <- DOWN
  12   <- DOWN
  13   <- DOWN
  14   <- DOWN
  15   <- DOWN

Y-FABRIC
   TO     0   1   2   3   4   5   6   7   8   9  10  11  12  13  14  15
FROM
  00      UP  UP  DN  DN  UNA UNA UNA UNA UNA UNA UNA UNA UNA UNA UNA UNA
  01      UP  UP  DN  DN  UNA UNA UNA UNA UNA UNA UNA UNA UNA UNA UNA UNA
  02      DN  DN  UP  UP  UNA UNA UNA UNA UNA UNA UNA UNA UNA UNA UNA UNA
  03      DN  DN  UP  UP  UNA UNA UNA UNA UNA UNA UNA UNA UNA UNA UNA UNA
  04   <- DOWN
  05   <- DOWN
  06   <- DOWN
  07   <- DOWN
  08   <- DOWN
  09   <- DOWN
  10   <- DOWN
  11   <- DOWN
  12   <- DOWN
  13   <- DOWN
  14   <- DOWN
  15   <- DOWN
```

Determine the Product Versions of the OSM Client Software

To determine the version of the OSM Low-Level Link:

Step	Action	Result
1.	Select **Start>Programs>HP OSM** and start OSM Low-Level Link.	The OSM Low-Level Link logon dialog box appears.
2.	Either log on or close the logon dialog box.	
3.	From the Help menu, Select **About HP OSM**.	The product version of the OSM Low-Level Link appears.
4.	Click **Close**.	The About OSM window disappears.

To determine the version of the OSM Notification Director:

Step	Action	Result
1.	Select **Start>Programs>HP OSM** and start OSM Notification Director.	The OSM Notification Director window appears.
2.	From the Help menu, Select **About HP OSM Notification Director**.	The product version of the OSM Notification Director appears.

Determine the Product Version of the TSM Client Software

Before you perform this operation, be sure that either the TSM Low-Level Link or the TSM Service Application is running.

Step	Action	Result
1.	In TSM, from the menu bar, click **Help**.	A menu appears.
2.	Select **About TSM**.	The product version of the TSM client software is listed.
3.	Click **Close**.	The About TSM window disappears.

Move the System Console

Step	Action
1.	Power off the system console.
2.	Use Step 19.2 Disconnect System Console on page 6-27 to disconnect the system console's Ethernet cables from the enclosure, block, or system you have stopped working with.
3.	Physically move the system console to its new location. Both the system console that you will connect to the enclosure and the enclosure itself must be offline (and therefore isolated). Until you configure the enclosure into a block, it cannot share a network with the target system. Refer to Step 7. Connect a System Console to the Enclosure on page 6-15.
4.	Connect the Ethernet cables to the enclosure, block, or system you now will be working with.
5.	Power on the system console.

When you log on to this enclosure, block, or system, you might receive warning messages about IP addresses. To resolve warning messages about IP addresses, see Step 11.2 Enter Enclosure IP Addresses on page 6-19.

Stop the OSM or TSM Low-Level Link

Step	Action	Result
1.	From the menu bar, select **File**.	A menu appears.
2.	Click **Exit.**	

Start a Startup TACL Session

Only one startup TACL session can be running at one time **on a system**. If another user on your system is already using the startup TACL connection, you receive a message that a duplicate session cannot be established. The startup TACL session is usually started on the system console.

Step	Action	Result
1.	From the OSM or TSM Low-Level Link menu bar, select **File**.	A menu appears.
2.	Select **Start Terminal Emulator.**	A menu appears.
3.	Select **For startup TACL**.	Two OutsideView windows appear.
4.	Look at both OutsideView windows. Because they are overlaid on each other, you need to switch between them. The active one displays a TACL prompt.	
5.	Log on as a super-group user.	

Start the OSM or TSM Low-Level Link

Step	Action	Result
1.	• In OSM, select **Start> Programs>HP OSM>OSM Low-Level Link**. • In TSM, select **Start> Programs>Compaq TSM>TSM Low-Level Link Application**.	A logon dialog box appears.
2.	Log on.	A menu appears.

Start the OSM Service Connection or TSM Service Application

Step	Action	Result
1.	• In OSM, select **Start> Programs>HP OSM>OSM Service Connection**. • In TSM, select **Start> Programs>Compaq TSM>TSM Service Application**.	A logon dialog box appears.
2.	Log on.	A menu appears.

B ServerNet Cabling

This appendix shows diagrams of all ServerNet cable connections. It also includes examples of how you can map these diagrams to actual NonStop S-series enclosures.

Example	Includes …
Correlation Between Cabling Diagram and One Enclosure with SEBS	• A sample ServerNet cabling diagram of one enclosure with connections shown to a SEB • The service side of a processor enclosure • How the diagram relates to the components in the enclosure • Table B-1, Router Port Connections to SEBs
Correlation Between Cabling Diagram and One Enclosure with MSEBS	• A sample ServerNet cabling diagram of one enclosure with connections shown to an MSEB • The service side of a processor enclosure • How the diagram relates to the components in the enclosure • Table B-2, Router Port Connections to MSEBs
Correlation Between Cabling Diagram and Two Enclosures	• A sample ServerNet cabling diagram of two connected enclosures • A processor enclosure connected to a NonStop S-series I/O enclosure • How the diagram relates to the components in each enclosure and the connections between them

Cabling Diagram

NonStop S-Series System — Tetra 8 Topology, X Fabric

NonStop S-Series System — Tetra 8 Topology, Y Fabric

NonStop S-Series System — Tetra 16 Topology, X Fabric

NonStop S-Series System — Tetra 16 Topology, Y Fabric

Note. The Tetra 16 cabling diagrams show the maximum possible number of I/O enclosures. Your service provider can tell you the number of I/O enclosures supported by the RVU running on your system.

Note. The diagrams in this guide identify all ServerNet expansion boards as SEBs. Your system might instead have or require modular ServerNet expansion boards (MSEBs) in the slots designated for SEBs. SEBs and MSEBs are functionally equivalent.

△ **Caution.** ServerNet cabling for IOAM enclosures is different from the cabling shown in this Appendix. Only service providers trained by HP should perform cabling of IOAM enclosures.

Example: Correlation Between Cabling Diagram and One Enclosure with SEBs

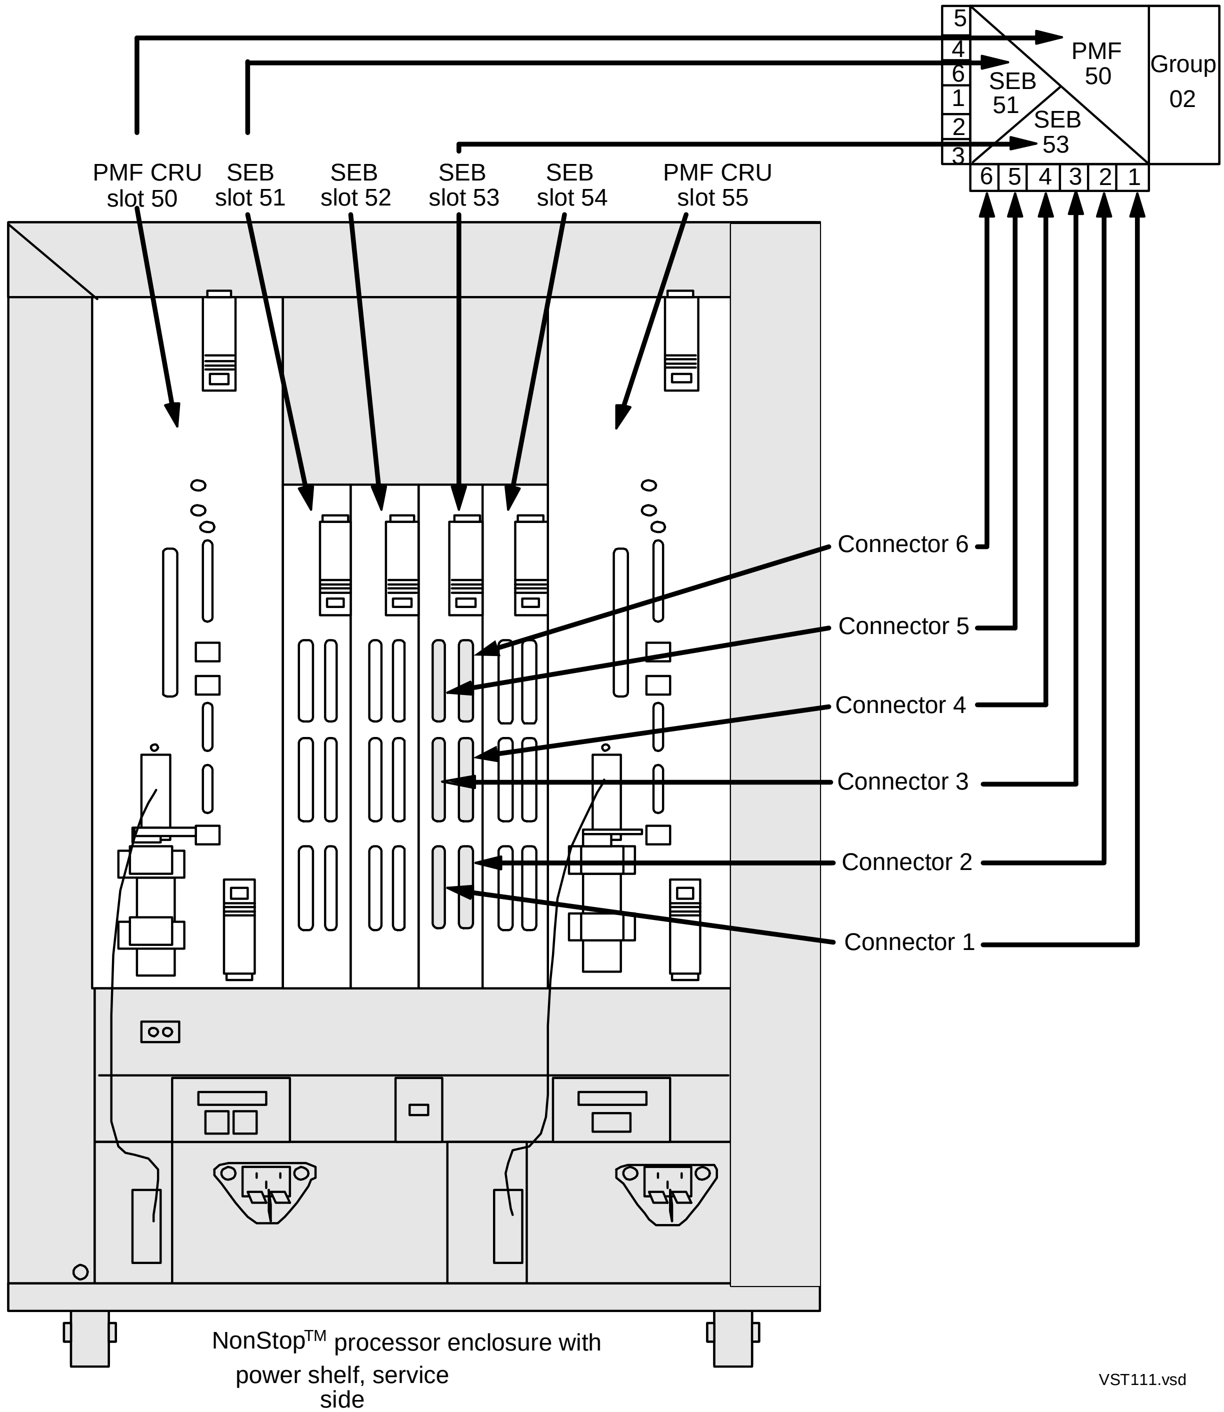

Table B-1. Router Port Connections to SEBs

Router Port	Connects toSEB
1	Cable connector 5
2	Cable connector 4
3	Cable connector 3
4	Cable connector 2
5	Cable connector 1

Example: **Correlation Between Cabling Diagram and One Enclosure with MSEBS**

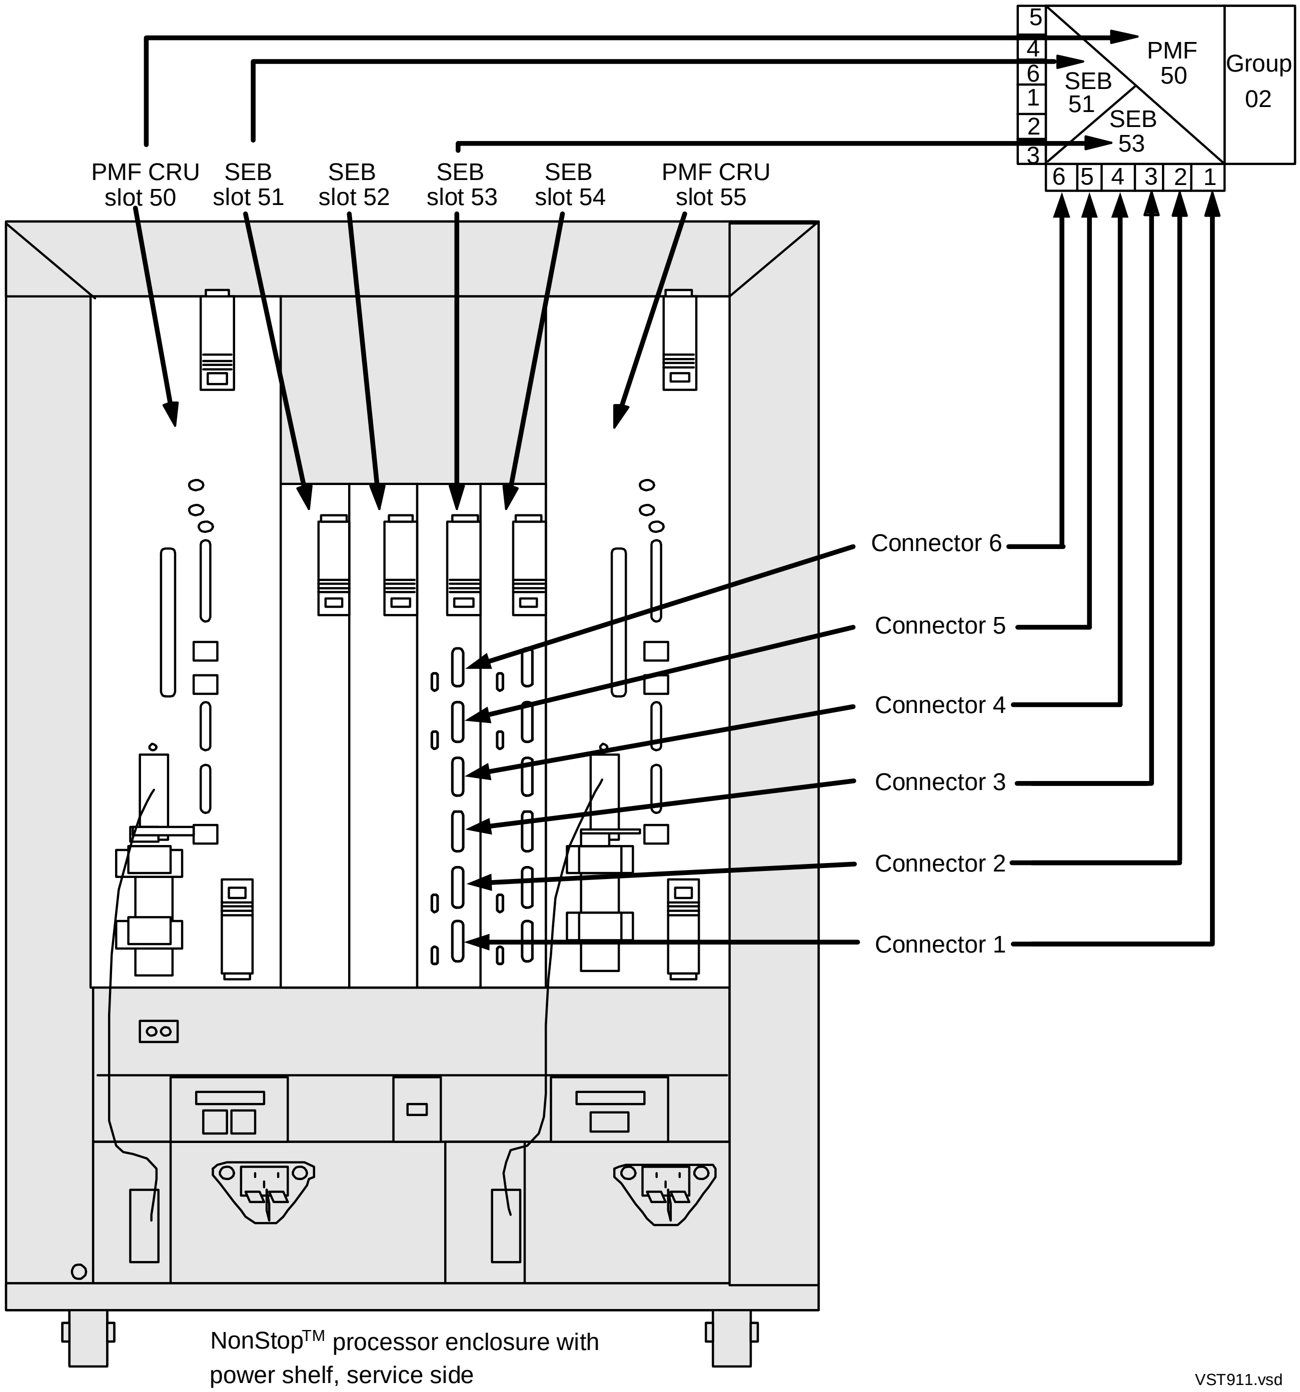

Table B-2. Router Port Connections to MSEBs

Router Port	Connects to MSEB
0	Cable connector 6
1	Cable connector 5
2	Cable connector 4
3	Cable connector 3
4	Cable connector 2
5	Cable connector 1

Example: Correlation Between Cabling Diagram and Two Enclosures

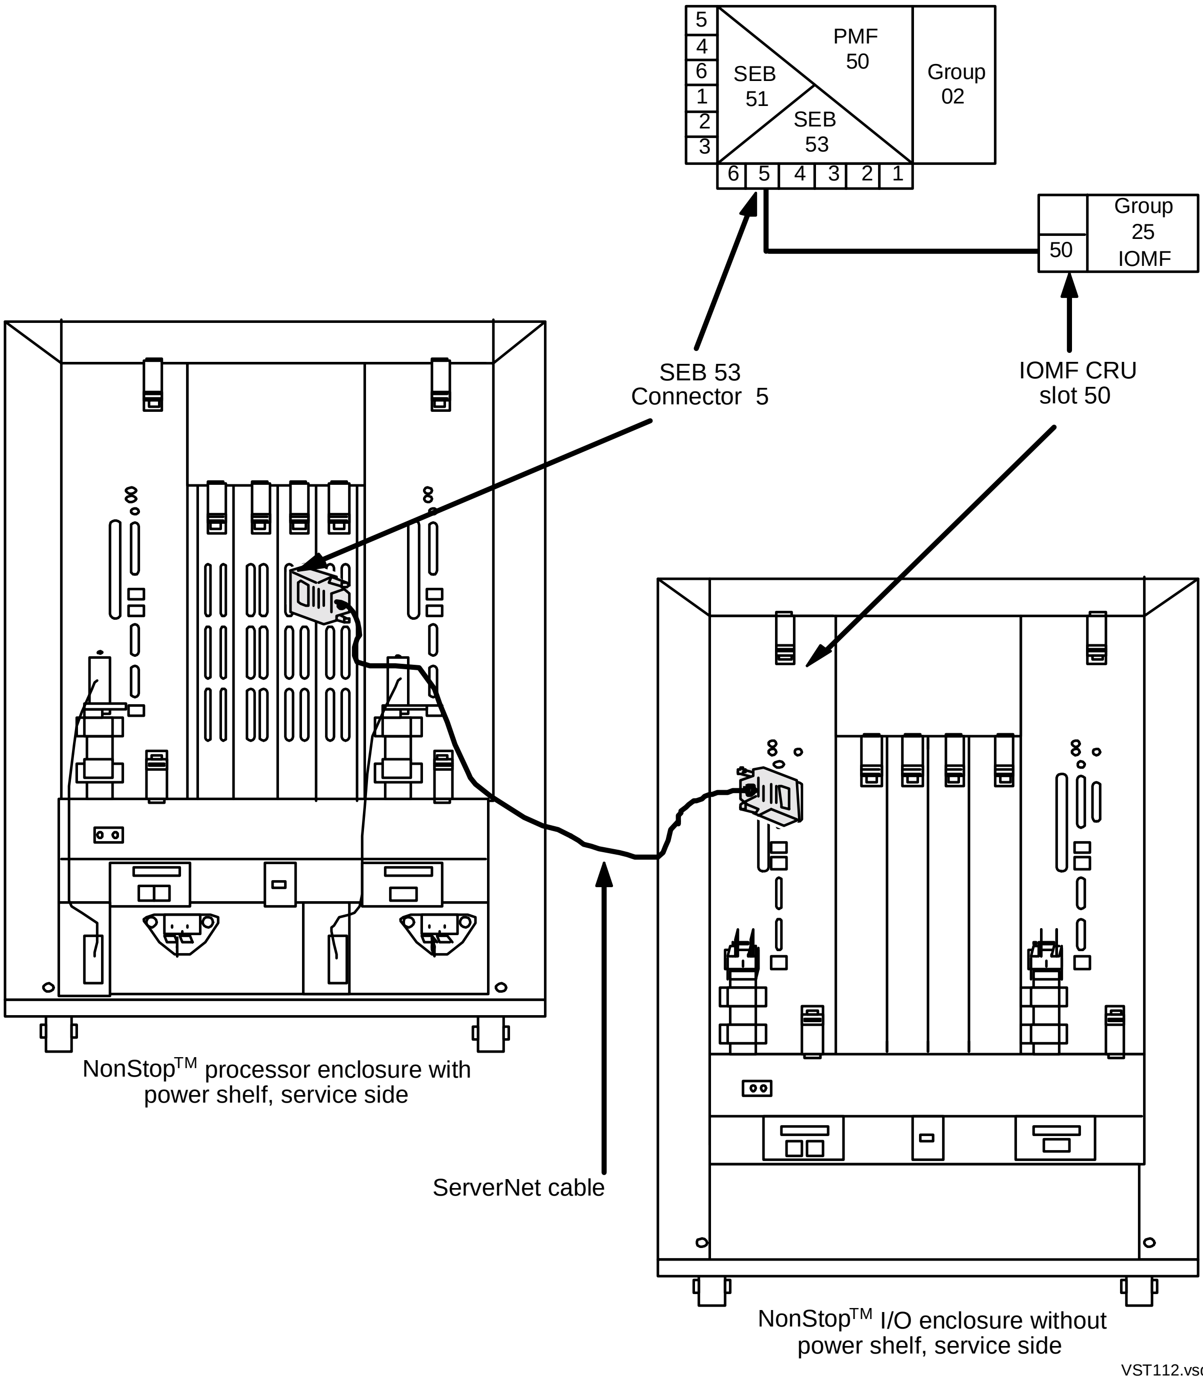

NonStop S-Series System — Tetra 8 Topology, X Fabric

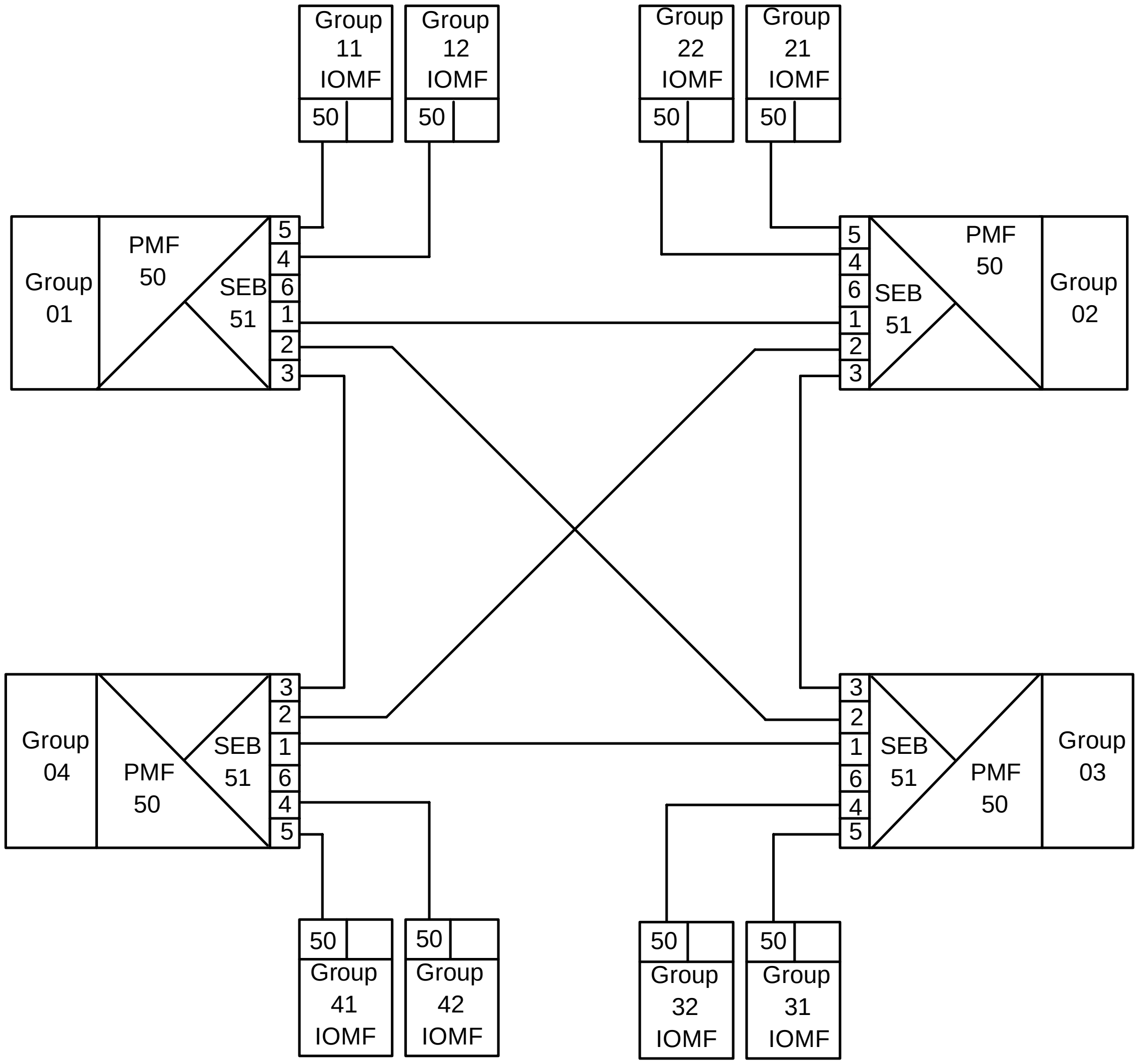

VST303.vsd

NonStop S-Series System — Tetra 8 Topology, Y Fabric

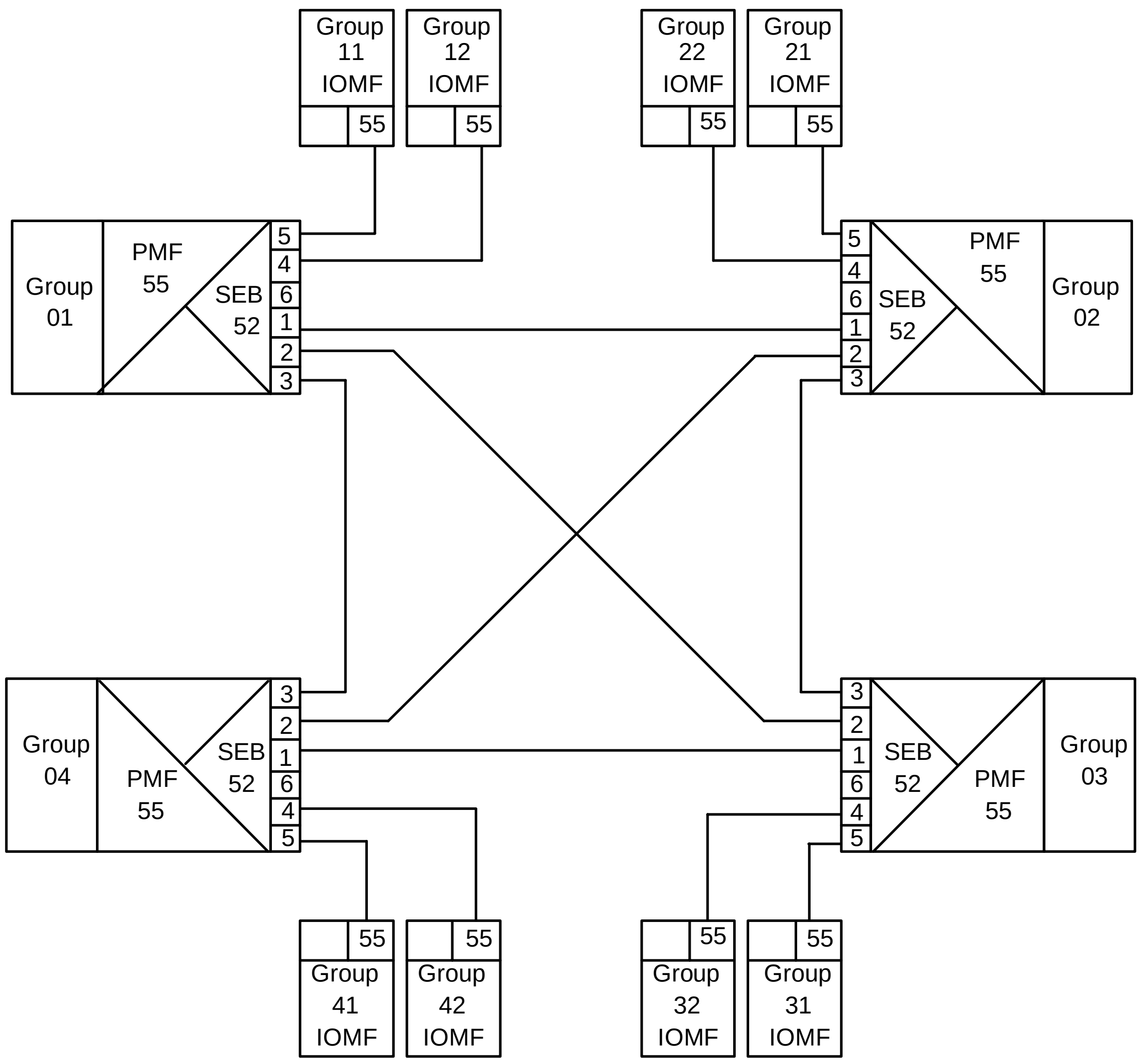

VST304.vsd

NonStop S-Series System — Tetra 16 Topology, X Fabric

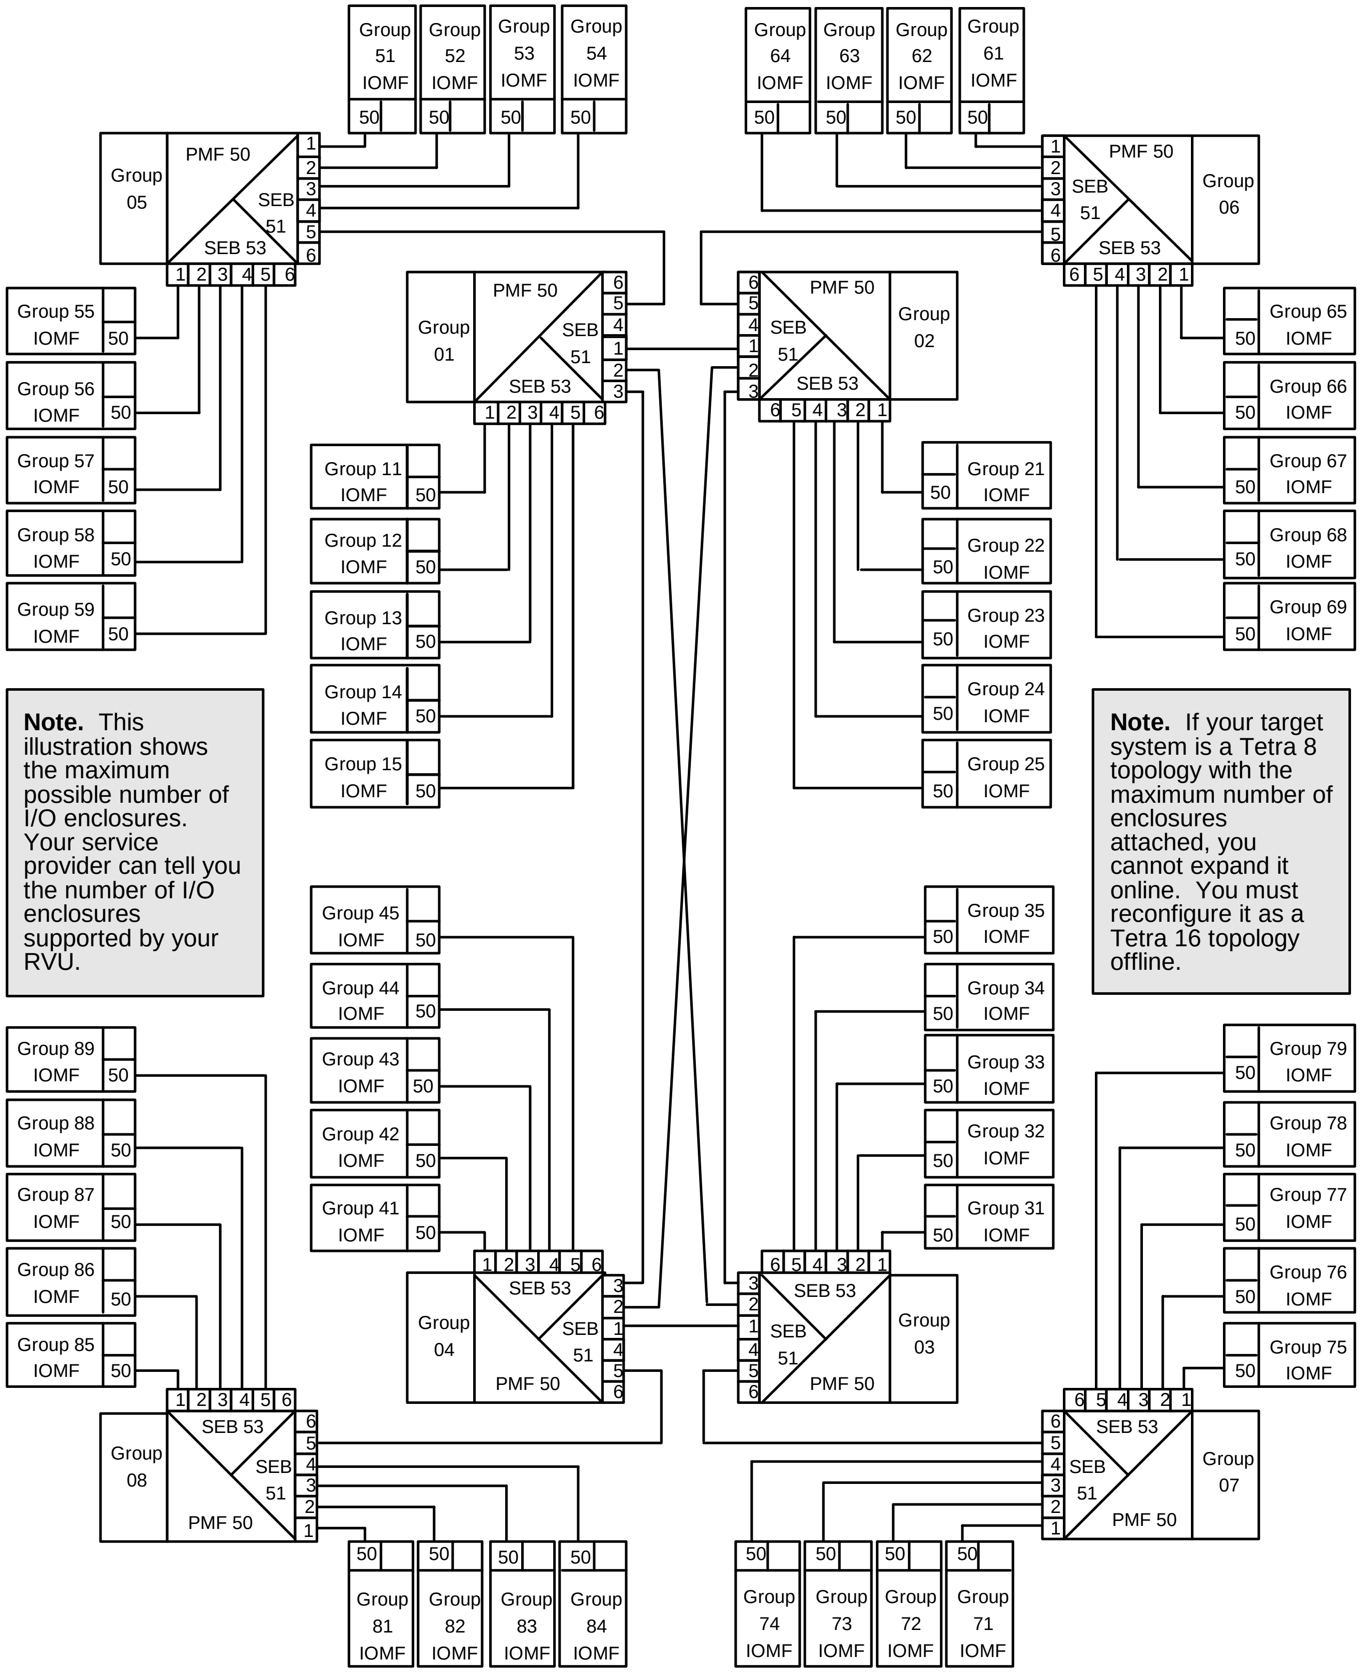

Note. This illustration shows the maximum possible number of I/O enclosures. Your service provider can tell you the number of I/O enclosures supported by your RVU.

Note. If your target system is a Tetra 8 topology with the maximum number of enclosures attached, you cannot expand it online. You must reconfigure it as a Tetra 16 topology offline.

VST301.vsd

NonStop S-Series System — Tetra 16 Topology, Y Fabric

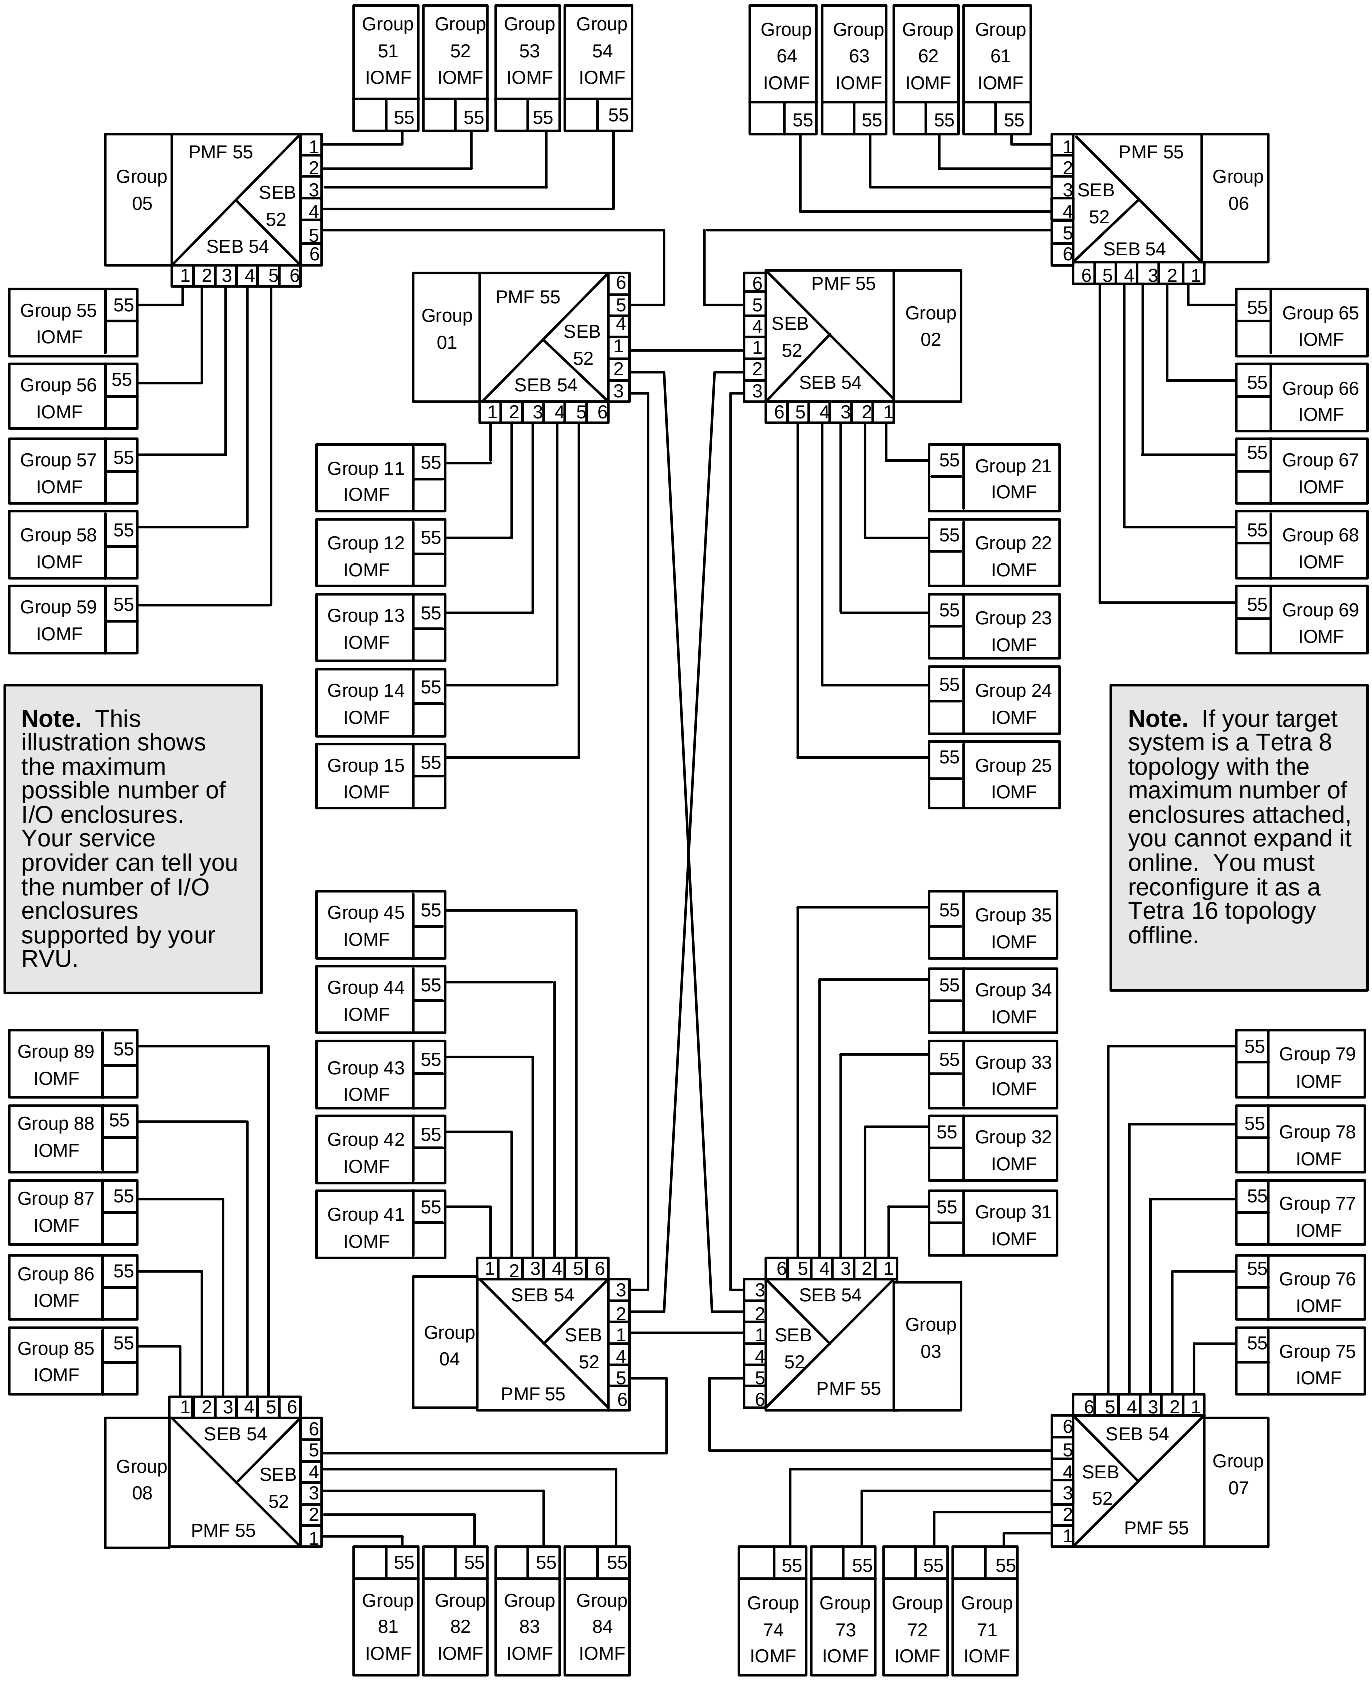

VST302.vsd

C Checklists and Worksheets

This appendix contains checklists and worksheets for planning system resizing.

Note. These worksheets show the maximum number of I/O enclosures. Your service provider can tell you the number of I/O enclosures supported by the RVU on your system.

Note. The diagrams in this guide identify all ServerNet expansion boards as SEBs. Your system might have or require modular ServerNet expansion boards (MSEBs) in the slots designated for SEBs. SEBs and MSEBs are functionally equivalent.

△ **Caution.** ServerNet cabling for IOAM enclosures is different from the cabling shown in the topology diagrams in this appendix. Only service providers trained by HP should perform cabling of IOAM enclosures.

Copy these worksheets and staple them together. You might need more than one copy of some worksheets. The checklists say which worksheets to copy more than once.

Sheet	Title	Purpose
Checklist	Tools and Supplies	Show tools and supplies for resizing
Checklist	Preliminary Procedures	Show how to prepare your site and system for resizing
Worksheet 1	OSM, TSM, System Information, and SP Firmware	Record OSM, TSM, service processor, and system information for target and donor systems (if applicable)
Worksheets 2.1 through 2.8 (2.*n*)	Groups 0*n* - *nn*	Record information about all groups in target and donor systems (if applicable)
Worksheet 2.1a	For IOAM Enclosures	Record information about IOAM enclosures in target and donor systems
Worksheet 3	Added Blocks — Tetra 8 Topology, X Fabric	Show on each fabric: • The configuration of your target system before resizing • Where to place added enclosures • What the group numbers will be for the added enclosures
Worksheet 4	Added Blocks — Tetra 8 Topology, Y Fabric	
Worksheet 5	Added Blocks — Tetra 16 Topology, X Fabric	
Worksheet 6	Added Blocks — Tetra 16 Topology, Y Fabric	
Worksheet 7	Block to Be Added	Record information about the block being added
Worksheet 8	Reduction: Hardware Inventory, Group # _____	Record information about components and processes in the enclosure to be removed
Worksheet 8a	Reduction: Hardware Inventory, IOAM Enclosure, Group # ____	Record information about components and processes in an IOAM enclosure to be removed

Sheet	Title	Purpose
Worksheet 9	ServerNet Worksheet: Tetra 8 Topology, X and Y Fabrics	Trace paths of ServerNet communication among enclosures
Worksheet 10	ServerNet Worksheet: Tetra 16 Topology, X Fabric	
Worksheet 11	ServerNet Worksheet: Tetra 16 Topology, Y Fabric	

Checklist: Tools and Supplies

- [ ] Short Phillips screwdriver 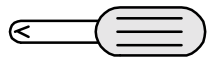
- [ ] Small slotted screwdriver 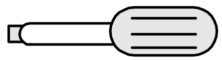
- [ ] 3/16” hex socket driver 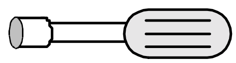
- [ ] 3/4” or 9/16” open-end wrench 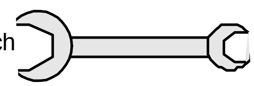
- [ ] Antistatic mat 
- [ ] Safety glasses 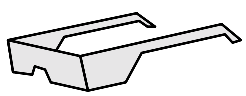
- [ ] Scissors or cutter 
- [ ] Tape 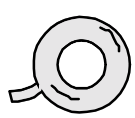
- [ ] Flashlight 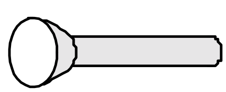
- [ ] Pointed object such as a straightened metal paper clip 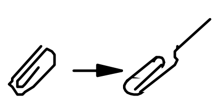
- [ ] NonStop™ S-series system console 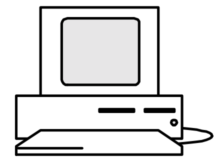
- [ ] NonStop™ S-series System Console Installer CD-ROM 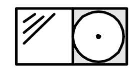
- [ ] Three (3) Ethernet cables 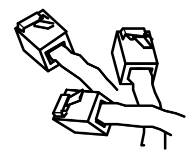
- [ ] One or more NonStop™ S-series system enclosures 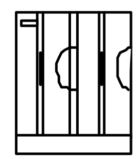
- [ ] Ethernet hub 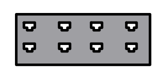
- [ ] SEBs, MSEBs, adapters, and other components you will need for your resized system 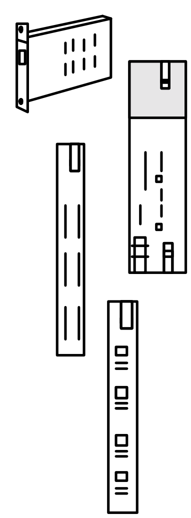

VST206.vsd

Checklist: Preliminary Procedures

- [ ] Ensure that all system consoles you will use to resize your system are running the appropriate version of OSM or TSM client software.

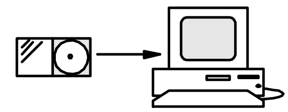

- [ ] If you are expanding your system, install enough MSEBs to accommodate the number of enclosures the system will contain after expansion.

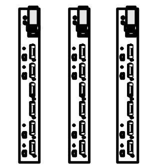

- [ ] Copy all worksheets and staple them together.

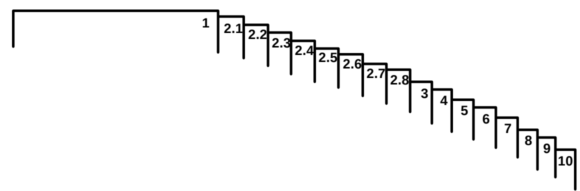

- [ ] To add more than three enclosures, make more copies of Worksheet 7: one extra for every two additional enclosures.

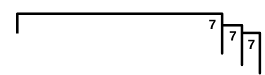

- [ ] To remove any enclosures, make more copies of Worksheet 8: one extra for every enclosure.

- [ ] To expand your system with enclosures removed from another system:
  - [ ] Make an extra copy of Worksheet 1, and circle *Donor System* at the top.
  - [ ] Make an extra copy of Worksheets 2.1 through 2.8, and circle *Donor System* at the top.
  - [ ] Staple the Donor System worksheets into a packet separate from the other worksheets.

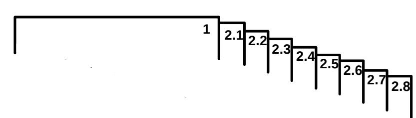

- [ ] If you are adding enclosures, roll them as close as possible towhere you will connect them.

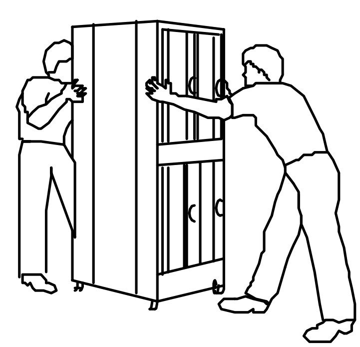

- [ ] Label all ServerNet cables with tape marked *X* or *Y* so that you do not cross-cable the fabrics.

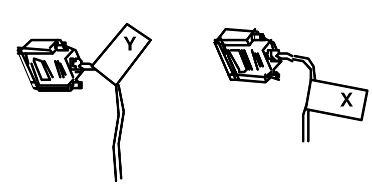

VST207.vsd

OSM, TSM, System, and Firmware Information

OSM Product Version

See the OSM User's Guide and the NonStop System Console Installer Guide for instructions on how to obtain this information.

TSM Product Version on the System Consoles

If your system consoles are not running the correct TSM client software for your RVU, upgrade the consoles with the appropriate NonStop System Console Installer CD-ROM.

TSM Client software on (circle one) **target** system console / **donor** system console

Initial Version: ______________________________

Needs to Be Upgraded to: ______________________________ Upgraded ☐

TSM Client software: block system console

Initial Version: ______________________________

Needs to Be Upgraded to: ______________________________ Upgraded ☐

System Information

System Name ______________________________

System Serial Number ______________________________

System Type ______________________________

Topology ☐ Tetra 8 ☐ Tetra 16

Group Numbers This system contains these group numbers before resizing (circle all that apply):

01 11 12 13 14 15	02 21 22 23 24 25	03 31 32 33 34 35	04 41 42 43 44 45
05 51 52 53 54	06 61 62 63 64	07 71 72 73 74	08 81 82 83 84
55 56 57 58 59	65 66 67 68 69	75 76 77 78 79	85 86 87 88 89

Current SYSnn ______________________________

Saved Configuration ______________________________

Saved Configuration
after Reduction (if applicable) ______________________________

SP Firmware Version ______________________________

For expansion: If the SP firmware of the block does not match the SP firmware of the target system, you must upgrade the block before continuing.

ME Firmware Version (IOAM only) ______________________________

ME FPGA Version (IOAM only) ______________________________

VST201.vsd

(circle one)**Target System / Donor System** Worksheet 2.1

Groups 01 11 12 13 14 15 (circle all that apply)

01	Slot **50**	Primary Service Connection IP Address		______
		Primary Low-Level Link IP Address		______
		Processor Firmware Version		______
		SCSI Controller Firmware Version		______
		Service Processor Firmware Version		______
	Slot **55**	Backup Service Connection IP Address		______
		Backup Low-Level Link IP Address		______
		Processor Firmware Version		______
		SCSI Controller Firmware Version		______
		Service Processor Firmware Version		______
11	Slot **50**	Low-Level Link IP Address	☐ Same as 01, slot 50 OR	______
		Service Processor Firmware	☐ Same as 01, slot 50 OR	______
	Slot **55**	Low-Level Link IP Address	☐ Same as 01, slot 55 OR	______
		Service Processor Firmware	☐ Same as 01, slot 55 OR	______
12	Slot **50**	Low-Level Link IP Address	☐ Same as 01, slot 50 OR	______
		Service Processor Firmware	☐ Same as 01, slot 50 OR	______
	Slot **55**	Low-Level Link IP Address	☐ Same as 01, slot 55 OR	______
		Service Processor Firmware	☐ Same as 01, slot 55 OR	
13	Slot **50**	Low-Level Link IP Address	☐ Same as 01, slot 50 OR	______
		Service Processor Firmware	☐ Same as 01, slot 50 OR	______
	Slot **55**	Low-Level Link IP Address	☐ Same as 01, slot 55 OR	______
		Service Processor Firmware	☐ Same as 01, slot 55 OR	______
14	Slot **50**	Low-Level Link IP Address	☐ Same as 01, slot 50 OR	______
		Service Processor Firmware	☐ Same as 01, slot 50 OR	______
	Slot **55**	Low-Level Link IP Address	☐ Same as 01, slot 55 OR	______
		Service Processor Firmware	☐ Same as 01, slot 55 OR	
15	Slot **50**	Low-Level Link IP Address	☐ Same as 01, slot 50 OR	______
		Service Processor Firmware	☐ Same as 01, slot 50 OR	______
	Slot **55**	Low-Level Link IP Address	☐ Same as 01, slot 55 OR	______
		Service Processor Firmware	☐ Same as 01, slot 55 OR	______

VST214.vsd

(circle one)**Target System / Donor System**

For IOAM Enclosures

Worksheet 2.1a

Groups ___ ___ ___ ___ ___ ___ (enter all that apply)

Processor Enclosure Group # ___

Slot **50**	Primary Service Connection IP Address	________
	Primary Low-Level Link IP Address	________
	Processor Firmware Version	________
	SCSI Controller Firmware Version	________
	Service Processor Firmware Version	________
Slot **55**	Backup Service Connection IP Address	________
	Backup Low-Level Link IP Address	________
	Processor Firmware Version	________
	SCSI Controller Firmware Version	________
	Service Processor Firmware Version	________

IOAM Enclosure Group # ___

Module 2	ME FPGA Version	________
	ME Firmware Version	________
Module 3	ME FPGA Version	________
	ME Firmware Version	________

IOAM Enclosure Group # ___

Module 2	ME FPGA Version	________
	ME Firmware Version	________
Module 3	ME FPGA Version	________
	ME Firmware Version	________

IOAM Enclosure Group # ___

Module 2	ME FPGA Version	________
	ME Firmware Version	________
Module 3	ME FPGA Version	________
	ME Firmware Version	________

IOAM Enclosure Group # ___

Module 2	ME FPGA Version	________
	ME Firmware Version	________
Module 3	ME FPGA Version	________
	ME Firmware Version	________

IOAM Enclosure Group # ___

Module 2	ME FPGA Version	________
	ME Firmware Version	________
Module 3	ME FPGA Version	________
	ME Firmware Version	________

VST224.vsd

(circle one)**Target System / Donor System** Worksheet 2.2

Groups 02 21 22 23 24 25 (circle all that apply)

02	Slot **50**	Primary Low-Level Link IP Address		______
		Processor Firmware Version		______
		SCSI Controller Firmware Version		______
		Service Processor Firmware Version		______
		Topology: ☐ Tetra 8 ☐ Tetra 16		
	Slot **55**	Backup Low-Level Link IP Address		______
		Processor Firmware Version		______
		SCSI Controller Firmware Version		______
		Service Processor Firmware Version		______
		Topology: ☐ Tetra 8 ☐ Tetra 16		
21	Slot **50**	Low-Level Link IP Address	☐ Same as 01, slot 50 OR	______
		Service Processor Firmware	☐ Same as 01, slot 50 OR	______
	Slot **55**	Low-Level Link IP Address	☐ Same as 01, slot 55 OR	______
		Service Processor Firmware	☐ Same as 01, slot 55 OR	______
22	Slot **50**	Low-Level Link IP Address	☐ Same as 01, slot 50 OR	______
		Service Processor Firmware	☐ Same as 01, slot 50 OR	______
	Slot **55**	Low-Level Link IP Address	☐ Same as 01, slot 55 OR	______
		Service Processor Firmware	☐ Same as 01, slot 55 OR	
23	Slot **50**	Low-Level Link IP Address	☐ Same as 01, slot 50 OR	______
		Service Processor Firmware	☐ Same as 01, slot 50 OR	______
	Slot **55**	Low-Level Link IP Address	☐ Same as 01, slot 55 OR	______
		Service Processor Firmware	☐ Same as 01, slot 55 OR	______
24	Slot **50**	Low-Level Link IP Address	☐ Same as 01, slot 50 OR	______
		Service Processor Firmware	☐ Same as 01, slot 50 OR	______
	Slot **55**	Low-Level Link IP Address	☐ Same as 01, slot 55 OR	______
		Service Processor Firmware	☐ Same as 01, slot 55 OR	
25	Slot **50**	Low-Level Link IP Address	☐ Same as 01, slot 50 OR	______
		Service Processor Firmware	☐ Same as 01, slot 50 OR	______
	Slot **55**	Low-Level Link IP Address	☐ Same as 01, slot 55 OR	______
		Service Processor Firmware	☐ Same as 01, slot 55 OR	______

VST215.vsd

Groups 03 31 32 33 34 35 (circle all that apply)

03	Slot **50**	Primary Low-Level Link IP Address			______
		Processor Firmware Version			______
		SCSI Controller Firmware Version			______
		Service Processor Firmware Version			______
		Topology: ☐ Tetra 8 ☐ Tetra 16			
	Slot **55**	Backup Low-Level Link IP Address			______
		Processor Firmware Version			______
		SCSI Controller Firmware Version			______
		Service Processor Firmware Version			______
		Topology: ☐ Tetra 8 ☐ Tetra 16			
31	Slot **50**	Low-Level Link IP Address	☐	Same as 01, slot 50 OR	______
		Service Processor Firmware	☐	Same as 01, slot 50 OR	______
	Slot **55**	Low-Level Link IP Address	☐	Same as 01, slot 55 OR	______
		Service Processor Firmware	☐	Same as 01, slot 55 OR	______
32	Slot **50**	Low-Level Link IP Address	☐	Same as 01, slot 50 OR	______
		Service Processor Firmware	☐	Same as 01, slot 50 OR	______
	Slot **55**	Low-Level Link IP Address	☐	Same as 01, slot 55 OR	______
		Service Processor Firmware	☐	Same as 01, slot 55 OR	
33	Slot **50**	Low-Level Link IP Address	☐	Same as 01, slot 50 OR	______
		Service Processor Firmware	☐	Same as 01, slot 50 OR	______
	Slot **55**	Low-Level Link IP Address	☐	Same as 01, slot 55 OR	______
		Service Processor Firmware	☐	Same as 01, slot 55 OR	______
34	Slot **50**	Low-Level Link IP Address	☐	Same as 01, slot 50 OR	______
		Service Processor Firmware	☐	Same as 01, slot 50 OR	______
	Slot **55**	Low-Level Link IP Address	☐	Same as 01, slot 55 OR	______
		Service Processor Firmware	☐	Same as 01, slot 55 OR	
35	Slot **50**	Low-Level Link IP Address	☐	Same as 01, slot 50 OR	______
		Service Processor Firmware	☐	Same as 01, slot 50 OR	______
	Slot **55**	Low-Level Link IP Address	☐	Same as 01, slot 55 OR	______
		Service Processor Firmware	☐	Same as 01, slot 55 OR	______

VST216.vsd

(circle one)**Target System / Donor System** Worksheet 2.4

Groups 04 41 42 43 44 45 (circle all that apply)

04 Slot **50** Primary Low-Level Link IP Address ____________________

Processor Firmware Version ____________________

SCSI Controller Firmware Version ____________________

Service Processor Firmware Version ____________________

Topology: ☐ Tetra 8 ☐ Tetra 16

Slot **55** Backup Low-Level Link IP Address ____________________

Processor Firmware Version ____________________

SCSI Controller Firmware Version ____________________

Service Processor Firmware Version ____________________

Topology: ☐ Tetra 8 ☐ Tetra 16

41 Slot **50** Low-Level Link IP Address ☐ Same as 01, slot 50 OR ____________________

Service Processor Firmware ☐ Same as 01, slot 50 OR ____________________

Slot **55** Low-Level Link IP Address ☐ Same as 01, slot 55 OR ____________________

Service Processor Firmware ☐ Same as 01, slot 55 OR ____________________

42 Slot **50** Low-Level Link IP Address ☐ Same as 01, slot 50 OR ____________________

Service Processor Firmware ☐ Same as 01, slot 50 OR ____________________

Slot **55** Low-Level Link IP Address ☐ Same as 01, slot 55 OR ____________________

Service Processor Firmware ☐ Same as 01, slot 55 OR

43 Slot **50** Low-Level Link IP Address ☐ Same as 01, slot 50 OR ____________________

Service Processor Firmware ☐ Same as 01, slot 50 OR ____________________

Slot **55** Low-Level Link IP Address ☐ Same as 01, slot 55 OR ____________________

Service Processor Firmware ☐ Same as 01, slot 55 OR ____________________

44 Slot **50** Low-Level Link IP Address ☐ Same as 01, slot 50 OR ____________________

Service Processor Firmware ☐ Same as 01, slot 50 OR ____________________

Slot **55** Low-Level Link IP Address ☐ Same as 01, slot 55 OR ____________________

Service Processor Firmware ☐ Same as 01, slot 55 OR

45 Slot **50** Low-Level Link IP Address ☐ Same as 01, slot 50 OR ____________________

Service Processor Firmware ☐ Same as 01, slot 50 OR ____________________

Slot **55** Low-Level Link IP Address ☐ Same as 01, slot 55 OR ____________________

Service Processor Firmware ☐ Same as 01, slot 55 OR ____________________

VST217.vsd

Groups 05 51 52 53 54 (circle all that apply)

05 Slot **50** Primary Low-Level Link IP Address ____________________

Processor Firmware Version ____________________

SCSI Controller Firmware Version ____________________

Service Processor Firmware Version ____________________

Topology: ☐ Tetra 8 ☐ Tetra 16

Slot **55** Backup Low-Level Link IP Address ____________________

Processor Firmware Version ____________________

SCSI Controller Firmware Version ____________________

Service Processor Firmware Version ____________________

Topology: ☐ Tetra 8 ☐ Tetra 16

51 Slot **50** Low-Level Link IP Address ☐ Same as 01, slot 50 OR ____________________

Service Processor Firmware ☐ Same as 01, slot 50 OR ____________________

Slot **55** Low-Level Link IP Address ☐ Same as 01, slot 55 OR ____________________

Service Processor Firmware ☐ Same as 01, slot 55 OR ____________________

52 Slot **50** Low-Level Link IP Address ☐ Same as 01, slot 50 OR ____________________

Service Processor Firmware ☐ Same as 01, slot 50 OR ____________________

Slot **55** Low-Level Link IP Address ☐ Same as 01, slot 55 OR ____________________

Service Processor Firmware ☐ Same as 01, slot 55 OR

53 Slot **50** Low-Level Link IP Address ☐ Same as 01, slot 50 OR ____________________

Service Processor Firmware ☐ Same as 01, slot 50 OR ____________________

Slot **55** Low-Level Link IP Address ☐ Same as 01, slot 55 OR ____________________

Service Processor Firmware ☐ Same as 01, slot 55 OR ____________________

54 Slot **50** Low-Level Link IP Address ☐ Same as 01, slot 50 OR ____________________

Service Processor Firmware ☐ Same as 01, slot 50 OR ____________________

Slot **55** Low-Level Link IP Address ☐ Same as 01, slot 55 OR ____________________

Service Processor Firmware ☐ Same as 01, slot 55 OR

VST218.vsd

Groups 06 61 62 63 64 (circle all that apply)

06	Slot **50**	Primary Low-Level Link IP Address		______
		Processor Firmware Version		______
		SCSI Controller Firmware Version		______
		Service Processor Firmware Version		______
		Topology: ☐ Tetra 8 ☐ Tetra 16		
	Slot **55**	Backup Low-Level Link IP Address		______
		Processor Firmware Version		______
		SCSI Controller Firmware Version		______
		Service Processor Firmware Version		______
		Topology: ☐ Tetra 8 ☐ Tetra 16		
61	Slot **50**	Low-Level Link IP Address	☐ Same as 01, slot 50 OR	______
		Service Processor Firmware	☐ Same as 01, slot 50 OR	______
	Slot **55**	Low-Level Link IP Address	☐ Same as 01, slot 55 OR	______
		Service Processor Firmware	☐ Same as 01, slot 55 OR	______
62	Slot **50**	Low-Level Link IP Address	☐ Same as 01, slot 50 OR	______
		Service Processor Firmware	☐ Same as 01, slot 50 OR	______
	Slot **55**	Low-Level Link IP Address	☐ Same as 01, slot 55 OR	______
		Service Processor Firmware	☐ Same as 01, slot 55 OR	
63	Slot **50**	Low-Level Link IP Address	☐ Same as 01, slot 50 OR	______
		Service Processor Firmware	☐ Same as 01, slot 50 OR	______
	Slot **55**	Low-Level Link IP Address	☐ Same as 01, slot 55 OR	______
		Service Processor Firmware	☐ Same as 01, slot 55 OR	______
64	Slot **50**	Low-Level Link IP Address	☐ Same as 01, slot 50 OR	______
		Service Processor Firmware	☐ Same as 01, slot 50 OR	______
	Slot **55**	Low-Level Link IP Address	☐ Same as 01, slot 55 OR	______
		Service Processor Firmware	☐ Same as 01, slot 55 OR	

VST218.vsd

Groups 07 71 72 73 74 (circle all that apply)

07 Slot **50** Primary Low-Level Link IP Address ____________

Processor Firmware Version ____________

SCSI Controller Firmware Version ____________

Service Processor Firmware Version ____________

Topology: ☐ Tetra 8 ☐ Tetra 16

Slot **55** Backup Low-Level Link IP Address ____________

Processor Firmware Version ____________

SCSI Controller Firmware Version ____________

Service Processor Firmware Version ____________

Topology: ☐ Tetra 8 ☐ Tetra 16

71	Slot **50**	Low-Level Link IP Address	☐ Same as 01, slot 50 OR	____________
		Service Processor Firmware	☐ Same as 01, slot 50 OR	____________
	Slot **55**	Low-Level Link IP Address	☐ Same as 01, slot 55 OR	____________
		Service Processor Firmware	☐ Same as 01, slot 55 OR	____________
72	Slot **50**	Low-Level Link IP Address	☐ Same as 01, slot 50 OR	____________
		Service Processor Firmware	☐ Same as 01, slot 50 OR	____________
	Slot **55**	Low-Level Link IP Address	☐ Same as 01, slot 55 OR	____________
		Service Processor Firmware	☐ Same as 01, slot 55 OR	
73	Slot **50**	Low-Level Link IP Address	☐ Same as 01, slot 50 OR	____________
		Service Processor Firmware	☐ Same as 01, slot 50 OR	____________
	Slot **55**	Low-Level Link IP Address	☐ Same as 01, slot 55 OR	____________
		Service Processor Firmware	☐ Same as 01, slot 55 OR	____________
74	Slot **50**	Low-Level Link IP Address	☐ Same as 01, slot 50 OR	____________
		Service Processor Firmware	☐ Same as 01, slot 50 OR	____________
	Slot **55**	Low-Level Link IP Address	☐ Same as 01, slot 55 OR	____________
		Service Processor Firmware	☐ Same as 01, slot 55 OR	

Groups 08 81 82 83 84 (circle all that apply)

08 Slot **50**
- Primary Low-Level Link IP Address ____________________
- Processor Firmware Version ____________________
- SCSI Controller Firmware Version ____________________
- Service Processor Firmware Version ____________________
- Topology: ☐ Tetra 8 ☐ Tetra 16

Slot **55**
- Backup Low-Level Link IP Address ____________________
- Processor Firmware Version ____________________
- SCSI Controller Firmware Version ____________________
- Service Processor Firmware Version ____________________
- Topology: ☐ Tetra 8 ☐ Tetra 16

81 Slot **50**
- Low-Level Link IP Address ☐ Same as 01, slot 50 OR ____________________
- Service Processor Firmware ☐ Same as 01, slot 50 OR ____________________

Slot **55**
- Low-Level Link IP Address ☐ Same as 01, slot 55 OR ____________________
- Service Processor Firmware ☐ Same as 01, slot 55 OR ____________________

82 Slot **50**
- Low-Level Link IP Address ☐ Same as 01, slot 50 OR ____________________
- Service Processor Firmware ☐ Same as 01, slot 50 OR ____________________

Slot **55**
- Low-Level Link IP Address ☐ Same as 01, slot 55 OR ____________________
- Service Processor Firmware ☐ Same as 01, slot 55 OR

83 Slot **50**
- Low-Level Link IP Address ☐ Same as 01, slot 50 OR ____________________
- Service Processor Firmware ☐ Same as 01, slot 50 OR ____________________

Slot **55**
- Low-Level Link IP Address ☐ Same as 01, slot 55 OR ____________________
- Service Processor Firmware ☐ Same as 01, slot 55 OR ____________________

84 Slot **50**
- Low-Level Link IP Address ☐ Same as 01, slot 50 OR ____________________
- Service Processor Firmware ☐ Same as 01, slot 50 OR ____________________

Slot **55**
- Low-Level Link IP Address ☐ Same as 01, slot 55 OR ____________________
- Service Processor Firmware ☐ Same as 01, slot 55 OR

Added Blocks: Tetra 8 Topology, X Fabric

- [ ] Shade in the enclosures, ports, and slot numbers that are in use in your current configuration. You can now see the available locations and group numbers for added enclosures.

- [ ] With a highlighter or a pen of a different color, circle where you will add the enclosures.

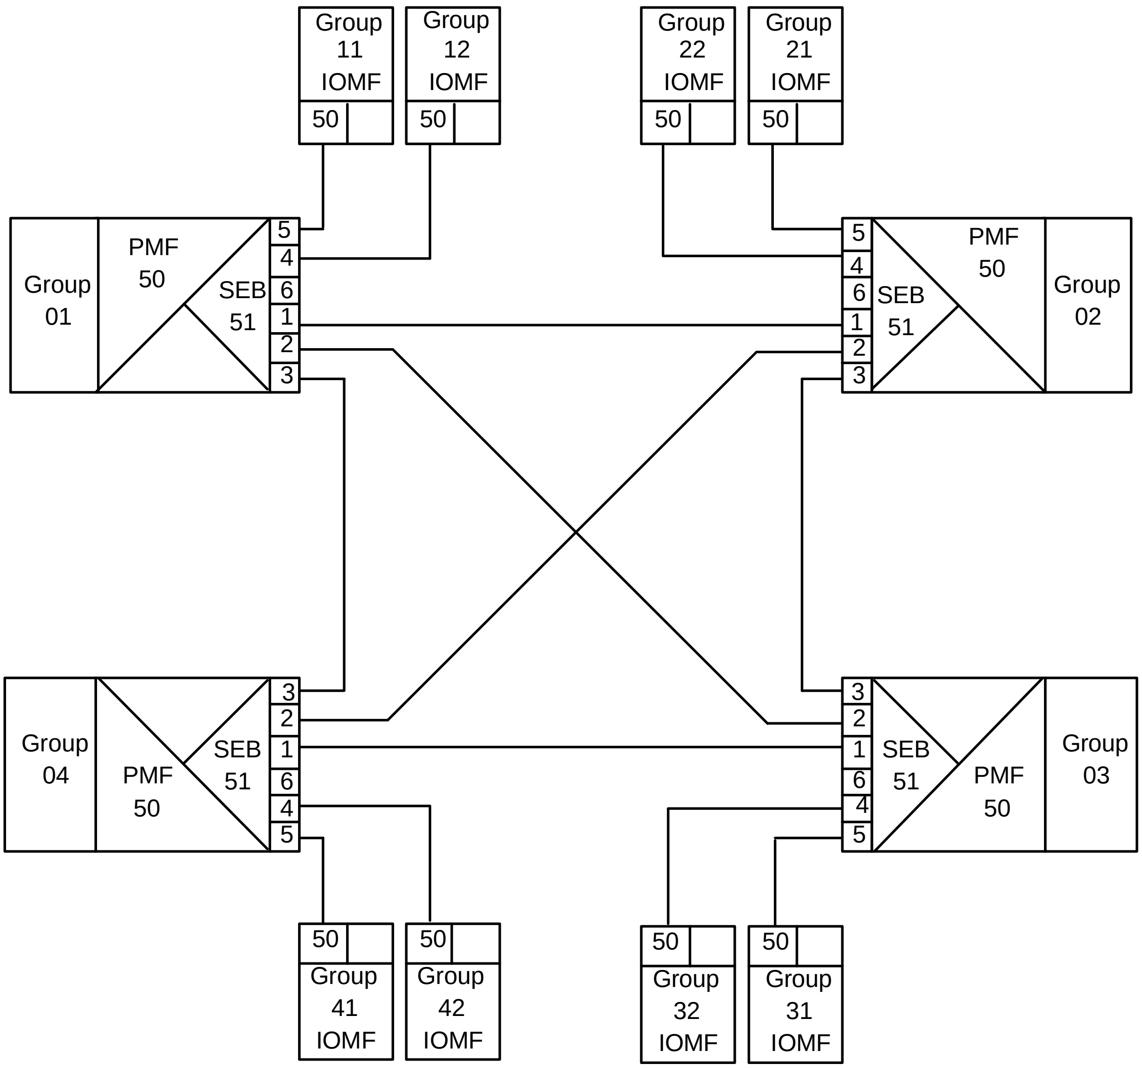

Added Blocks: Tetra 8 Topology, Y Fabric

☐ Shade in the enclosures, ports, and slot numbers that are in use in your current configuration. You can now see the available locations and group numbers for added enclosures.

☐ With a highlighter or a pen of a different color, circle where you will add the enclosures.

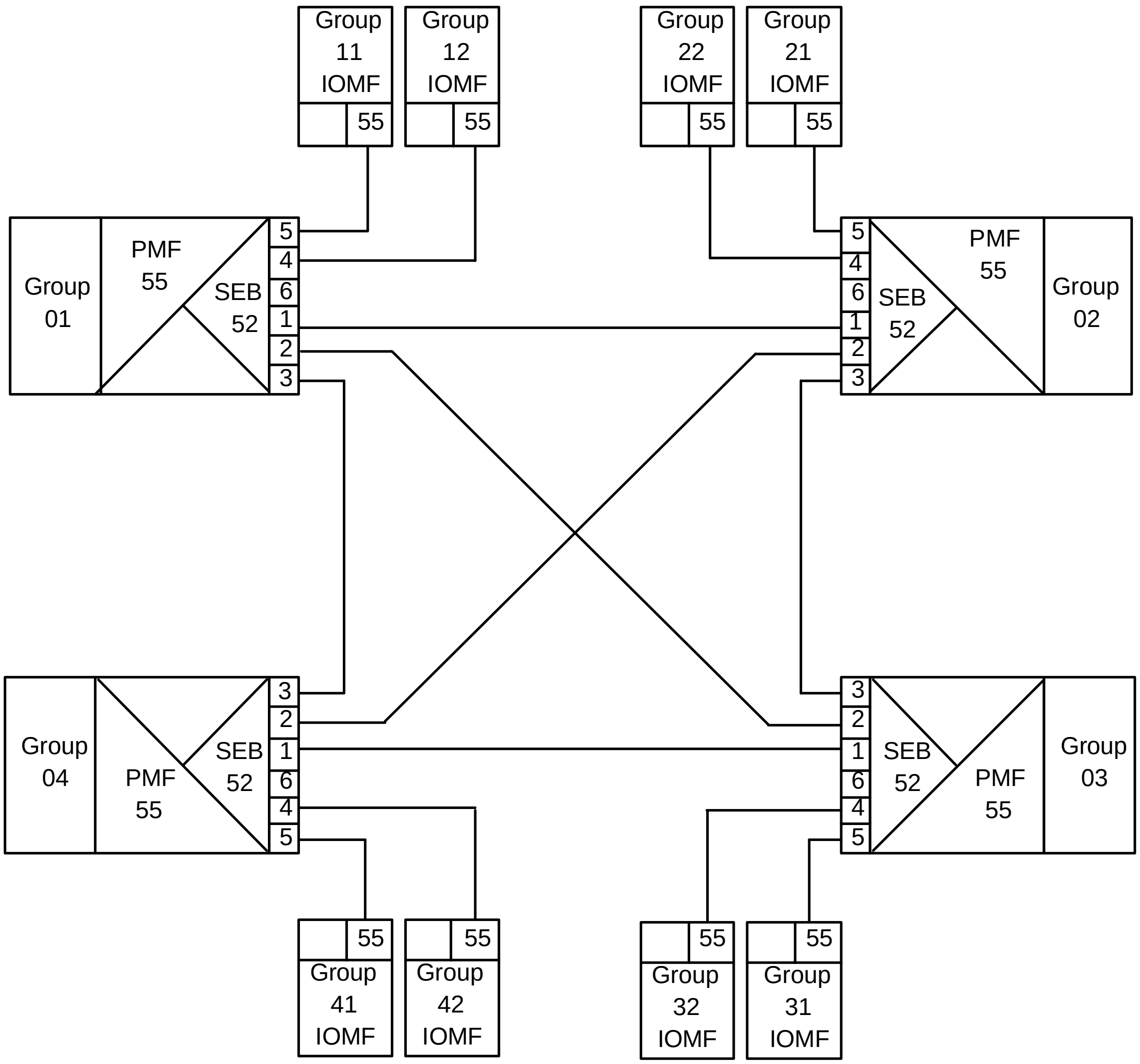

Added Blocks: Tetra 16 Topology, X Fabric

☐ Shade in the enclosures, ports, and slot numbers that are in use in your current configuration. You can now see the available locations and group numbers for added enclosures.

☐ With a highlighter or a pen of a different color, circle where you will add the enclosures.

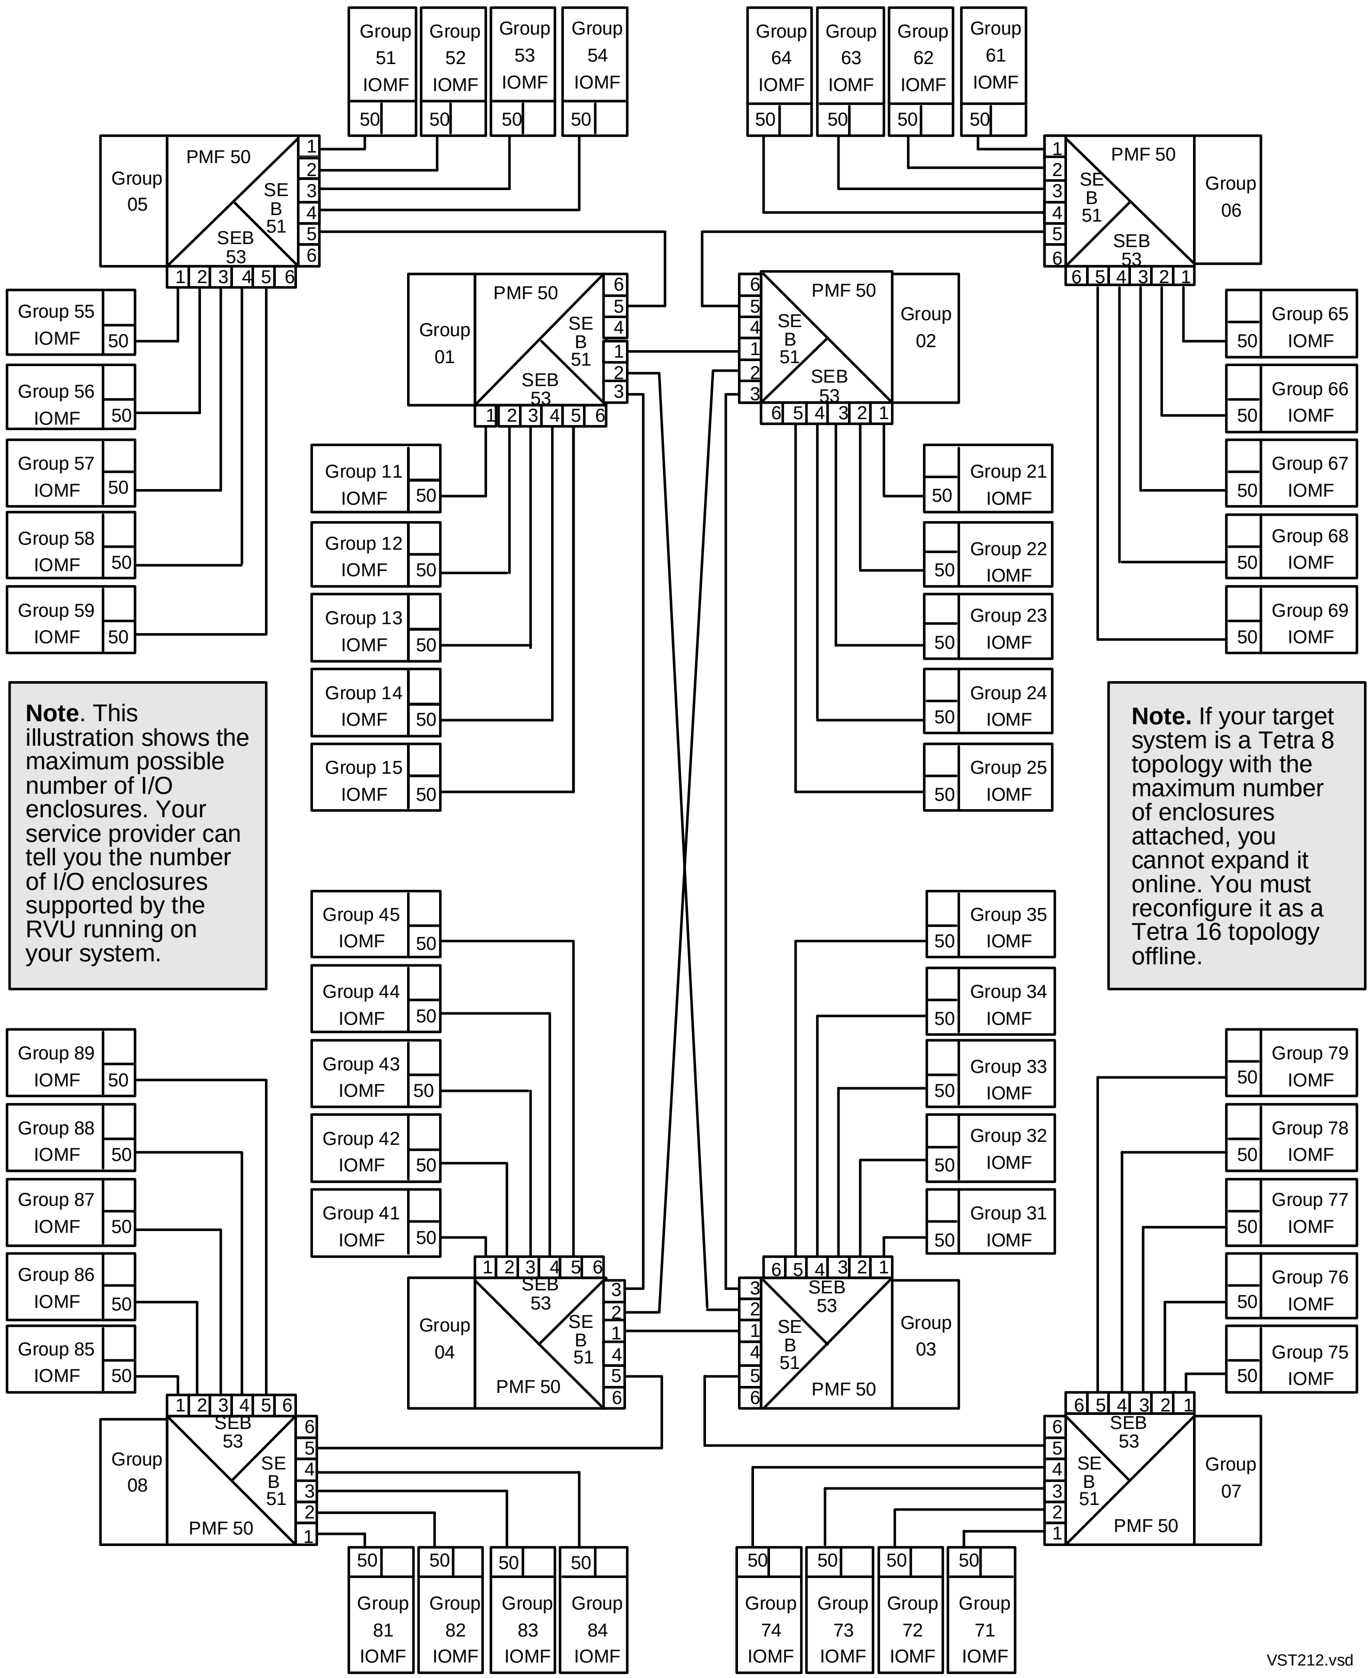

Note. This illustration shows the maximum possible number of I/O enclosures. Your service provider can tell you the number of I/O enclosures supported by the RVU running on your system.

Note. If your target system is a Tetra 8 topology with the maximum number of enclosures attached, you cannot expand it online. You must reconfigure it as a Tetra 16 topology offline.

VST212.vsd

Added Blocks: Tetra 16 Topology, Y Fabric

☐ Shade in the enclosures, ports, and slot numbers that are in use in your current configuration. You can now see the available locations and group numbers for added enclosures.

☐ With a highlighter or a pen of a different color, circle where you will add the enclosures.

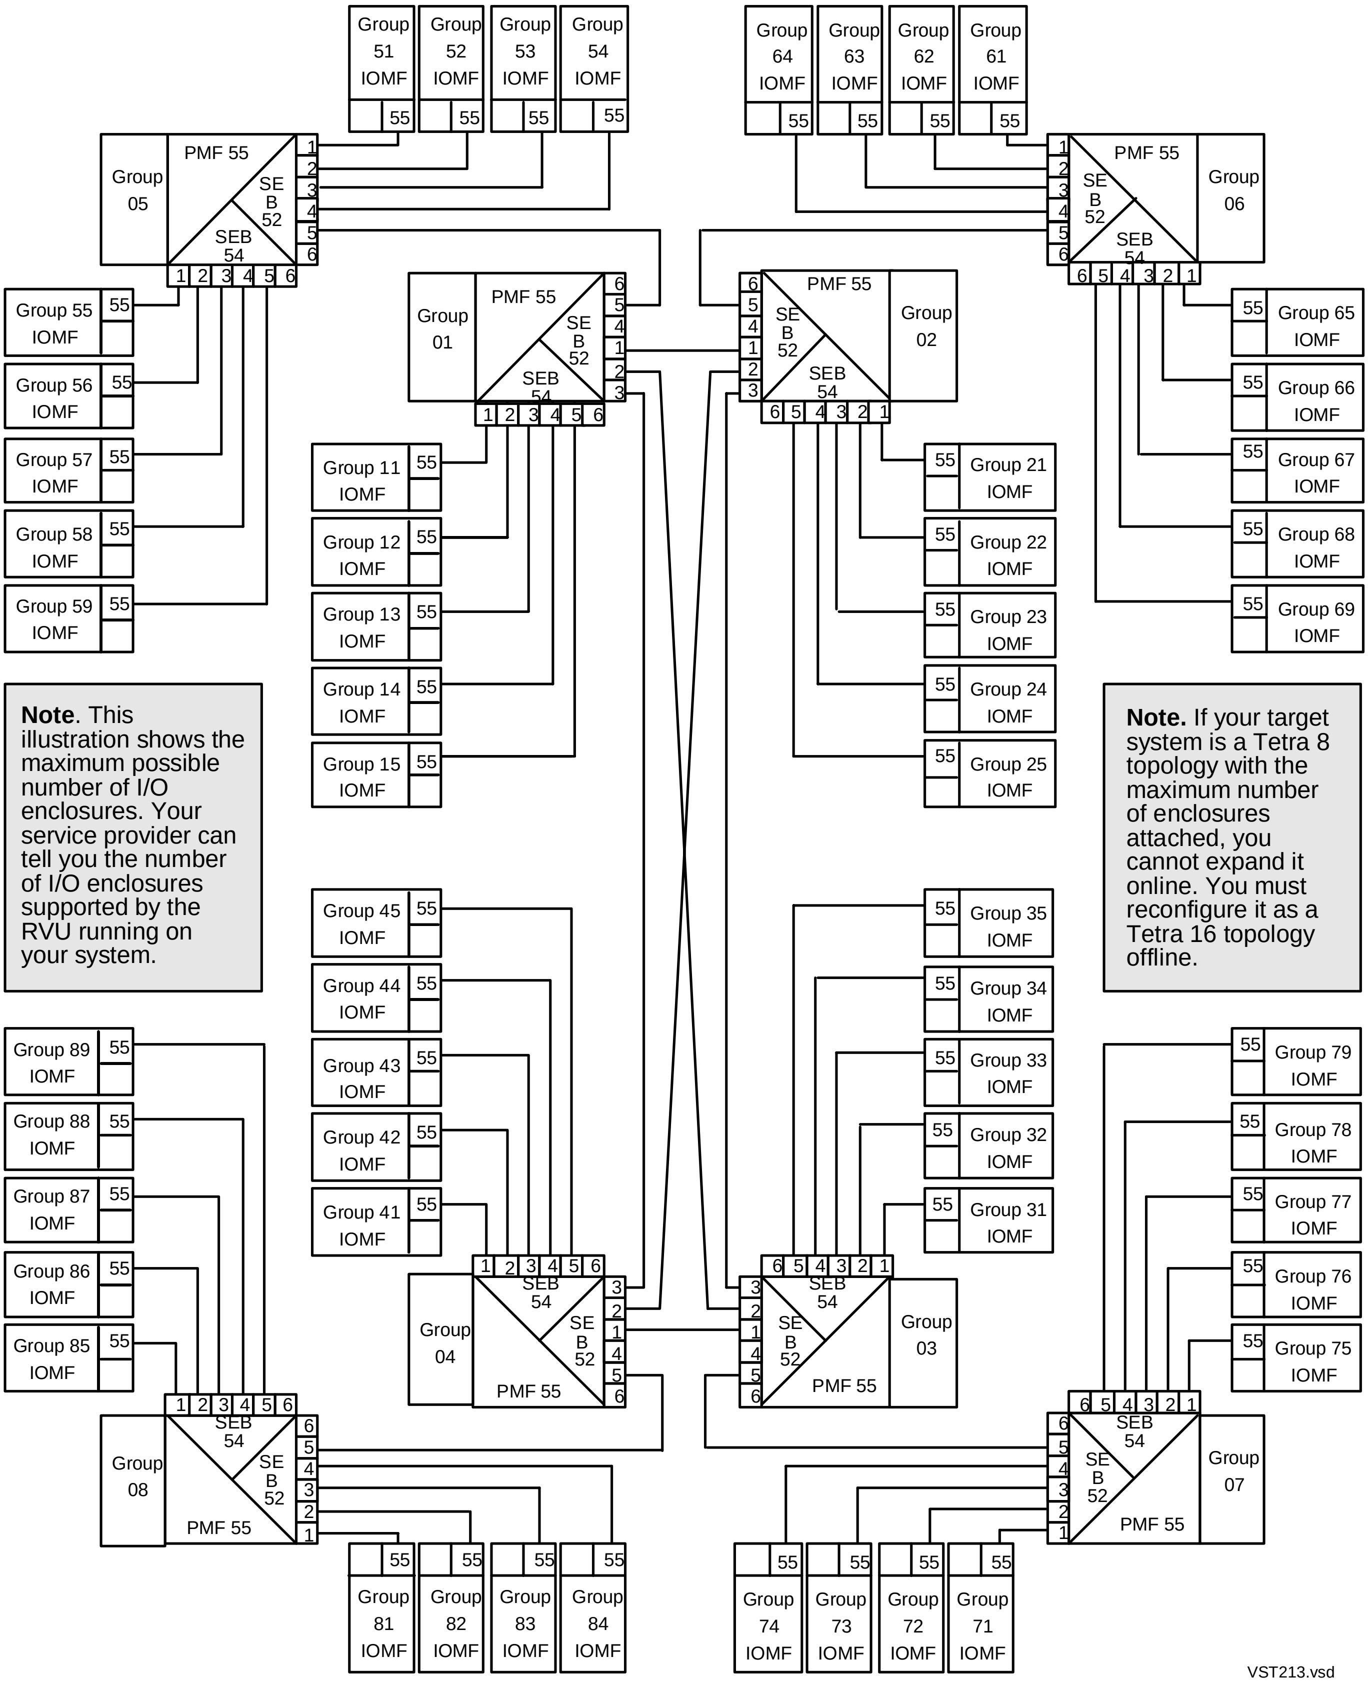

Note. This illustration shows the maximum possible number of I/O enclosures. Your service provider can tell you the number of I/O enclosures supported by the RVU running on your system.

Note. If your target system is a Tetra 8 topology with the maximum number of enclosures attached, you cannot expand it online. You must reconfigure it as a Tetra 16 topology offline.

VST213.vsd

Block to Be Added

☐ If your block contains more enclosures than you can record on this sheet, copy this sheet and continue recording on that copy. Fill in the “1 of ____ pages” blank to the right.

☐ If you are adding a block made of enclosures from another system, do not use this sheet. Copy Worksheets 2.1 through 2.8 and circle “Donor System” at the top instead.

Target System Information

☐ Tetra 8 ☐ Tetra 16 SP Firmware Version ____________________

Processor Enclosure: to be added to target system as Group __________

Slot 50

Low-Level Link IP Address	192.231.36.2 OR	____________________
Subnet Mask	255.255.255.0 OR	____________________
Gateway Route	192.231.36.9 OR	____________________
SP Firmware Version		____________________
		☐ same as target system ☐ upgraded

Slot 55

Low-Level Link IP Address	192.231.36.3 OR	____________________
Subnet Mask	255.255.255.0 OR	____________________
Gateway Route	192.231.36.9 OR	____________________
SP Firmware Version		____________________
		☐ same as target system ☐ upgraded

I/O Enclosure: to be added to target system as Group __________

Slot 50

Low-Level Link IP Address	192.231.36.2 OR	____________________
SP Firmware Version		____________________
		☐ same as target system ☐ upgraded

Slot 55

Low-Level Link IP Address	192.231.36.3 OR	____________________
SP Firmware Version		____________________
		☐ same as target system ☐ upgraded

I/O Enclosure: to be added to target system as Group __________

Slot 50

Low-Level Link IP Address	192.231.36.2 OR	____________________
SP Firmware Version		____________________
		☐ same as target system ☐ upgraded

Slot 55

Low-Level Link IP Address	192.231.36.3 OR	____________________
SP Firmware Version		____________________
		☐ same as target system ☐ upgraded

IOAM Enclosure: to be added to target system as Group __________

Module 2	ME Firmware Version	____________________
	ME FPGA Version	____________________
Module 3	ME Firmware Version	____________________
	ME FPGA Version	____________________

VST202.vsd

Reduction: Hardware Inventory, Group #__________

Device Type	Device Name	Slot
DISK	____________________	1
DISK	____________________	2
DISK	____________________	3
DISK	____________________	4
DISK	____________________	5
DISK	____________________	6
DISK	____________________	7
DISK	____________________	8
DISK	____________________	11
DISK	____________________	12
DISK	____________________	13
DISK	____________________	14
DISK	____________________	15
DISK	____________________	16
DISK	____________________	17
DISK	____________________	18

Device Type (circle one)	Device Name	Slot
SEB MSEB /FX E4SA 6760 Other SWAN ∇__________	____________________	51
SEB MSEB /FX E4SA 6760 Other SWAN ∇__________	____________________	52
SEB MSEB /FX E4SA 6760 Other SWAN ∇__________	____________________	53
SEB MSEB /FX E4SA 6760 Other SWAN ∇__________	____________________	54
TAPE DRIVE	____________________	50
TAPE DRIVE	____________________	55

Critical processes to move to other enclosures
(Use additional pieces of paper as necessary.)

VST209.vsd

Reduction: Hardware Inventory, IOAM Enclosure, Group #__________

Module 2

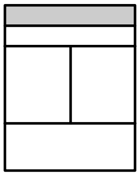

Device Type | Device Name | Slot
ServerNet Switch Board ________________________ 14

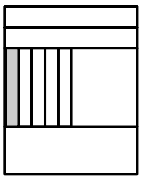

Device Type (circle one) | Device Name | Slot
FCSA G4SA ________________________ 1

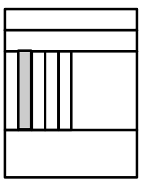

Device Type (circle one) | Device Name | Slot
FCSA G4SA ________________________ 2

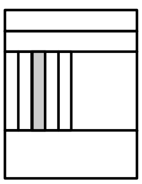

Device Type (circle one) | Device Name | Slot
FCSA G4SA ________________________ 3

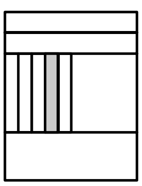

Device Type (circle one) | Device Name | Slot
FCSA G4SA ________________________ 4

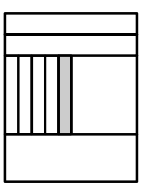

Device Type (circle one) | Device Name | Slot
FCSA G4SA ________________________ 5

Module 3

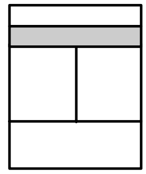

Device Type | Device Name | Slot
ServerNet Switch Board ________________________ 14

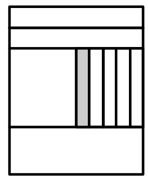

Device Type (circle one) | Device Name | Slot
FCSA G4SA ________________________ 1

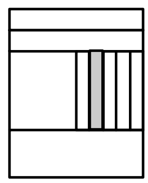

Device Type (circle one) | Device Name | Slot
FCSA G4SA ________________________ 2

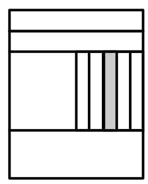

Device Type (circle one) | Device Name | Slot
FCSA G4SA ________________________ 3

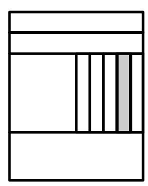

Device Type (circle one) | Device Name | Slot
FCSA G4SA ________________________ 4

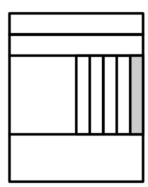

Device Type (circle one) | Device Name | Slot
FCSA G4SA ________________________ 5

Critical processes to move to other enclosures (Use additional pieces of paper as necessary.)

VST208.vsd

ServerNet Worksheet: Tetra 8 Topology, X and Y Fabrics

- ☐ With a highlighter, shade the enclosures, ports, and slot numbers that are in use in your current configuration.
- ☐ With a pen, black out the enclosure you plan to remove.
- ☐ Starting from the central “X” in this figure (representing the tetrahedral topology for the system), trace all ServerNet cables to the ServerNet expansion boards (SEBs) or modular SEBs (MSEBs) in the processor enclosures, and trace through all those SEBs or MSEBs to the I/O enclosures attached to the processor enclosures. However, when you trace pathways to the remaining enclosures in your system, do not trace through the blacked-out area. All pathways end when they reach that enclosure.
- ☐ Circle all enclosures that are affected by the enclosure removal. List them in reverse numeric order, highest to lowest. This is the order in which you will remove them from your system.

X Fabric

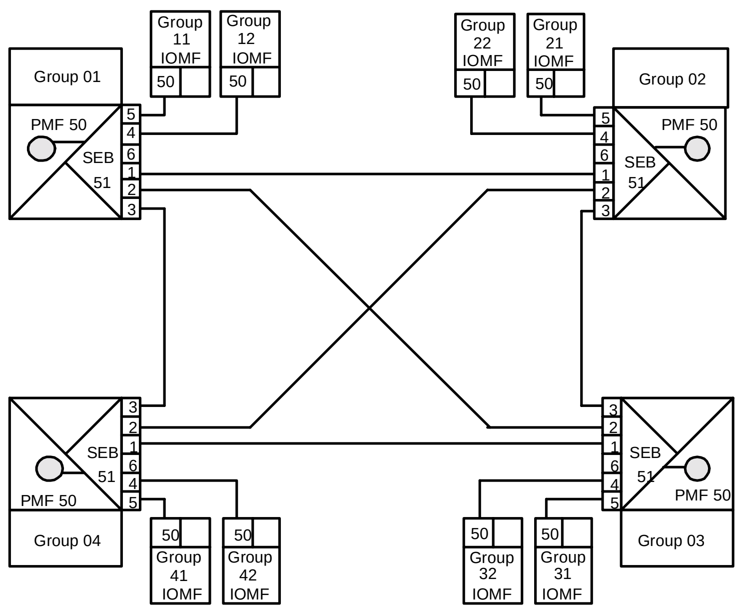

Note1 : This illustration shows the maximum possible number of I/O enclosures. Your service provider can tell you the number of I/O enclosures supported by your RVU.

Note 2: The router connections still exist between the PMF CRU in slot 50 of each enclosure and the processor in the PMF CRU in slot 55 of each enclosure. However, in this figure they are not shown.

Y Fabric

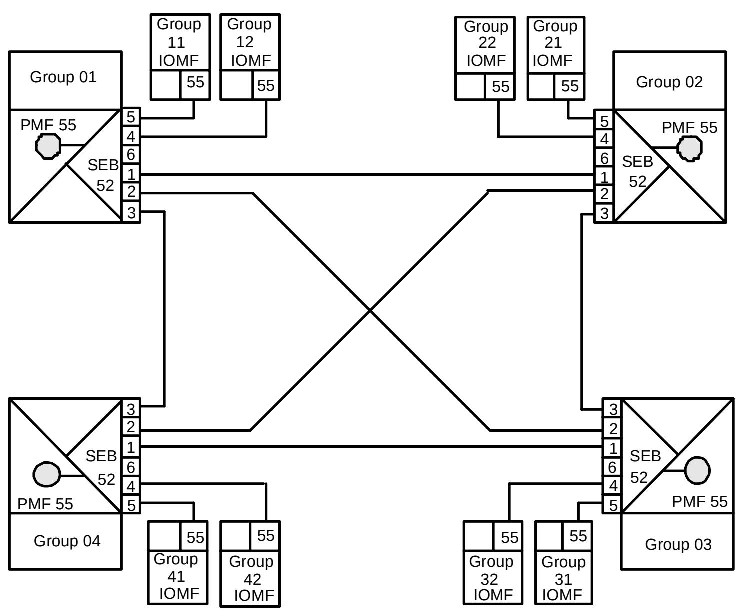

Enclosures to be Removed
(list in reverse numerical order)

VST386.vsd

ServerNet Communication Pathways: Tetra 16 Topology, X Fabric

- ☐ With a highlighter, shade the enclosures, ports, and slot numbers that are in use in your current configuration.
- ☐ With a pen, black out the CRU you plan to remove.
- ☐ If you are removing a PMF CRU, check this box. Then, on the worksheet for the Y fabric, black out the PMF CRU in slot 55 in the corresponding enclosure to indicate that this PMF CRU will lose access to the X fabric as well.
- ☐ On this worksheet, trace pathways of communication through all routers and ServerNet cables that connect to the CRU you plan to remove. Circle all components and enclosures that will be affected by that CRU removal.

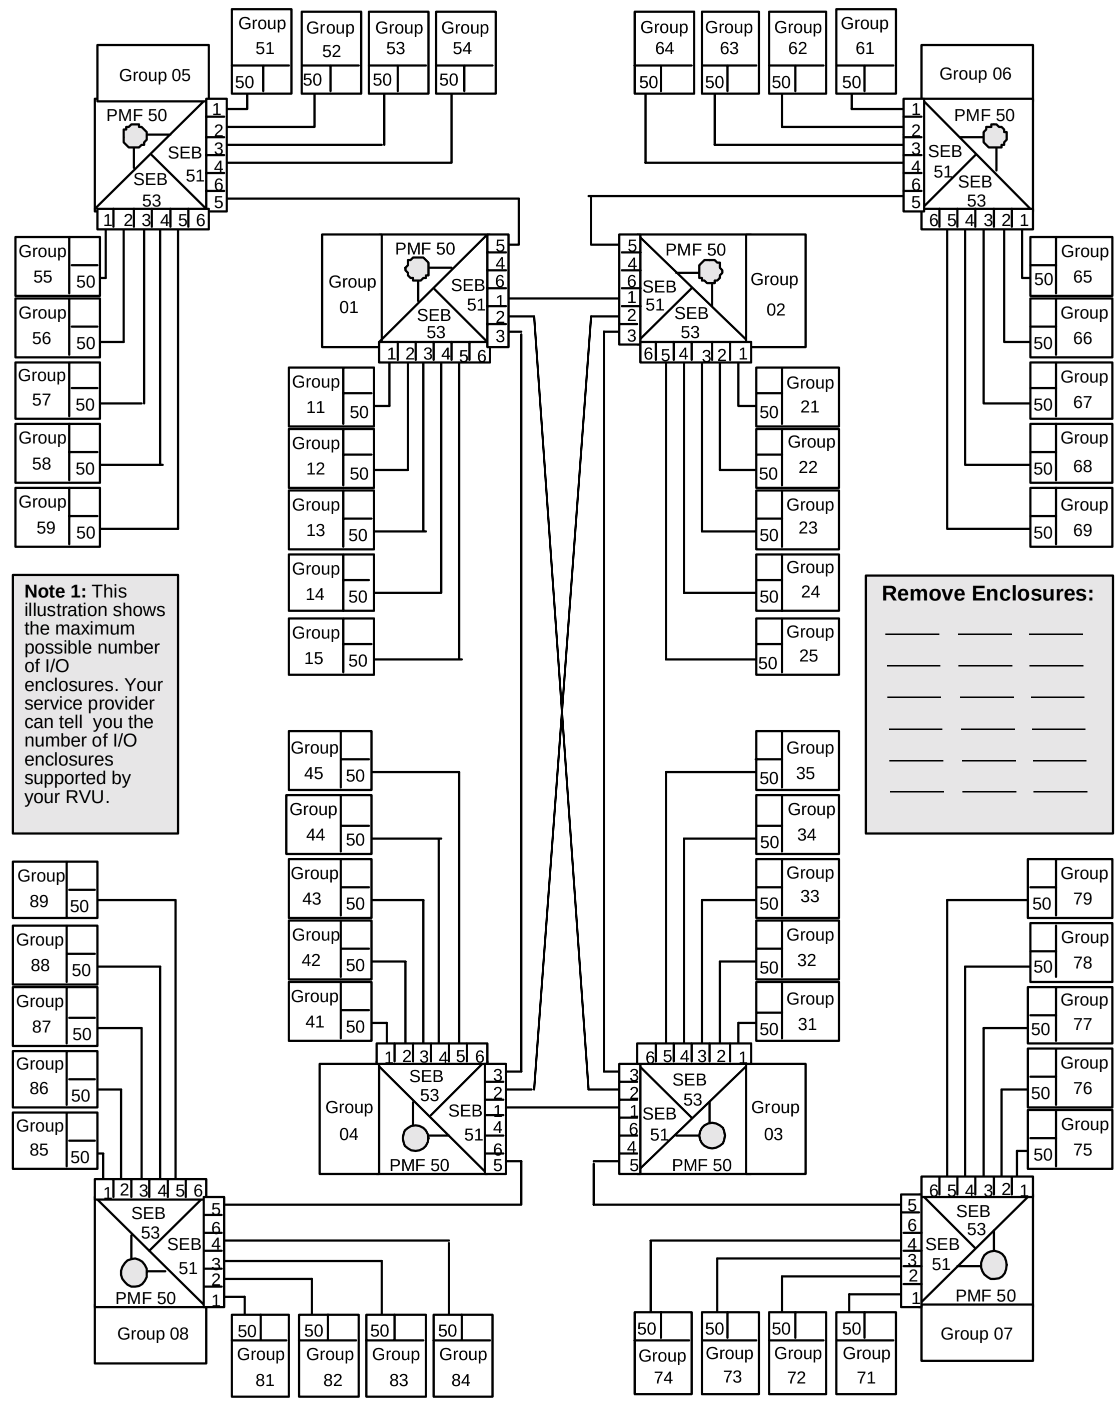

ServerNet Communication Pathways: Tetra 16 Topology, Y Fabric

- [ ] With a highlighter, shade the enclosures, ports, and slot numbers that are in use in your current configuration.
- [ ] With a pen, black out the CRU you plan to remove.
- [ ] If you are removing a PMF CRU, check this box. Then, on the worksheet for the X fabric, black out the PMF CRU in slot 50 in the corresponding enclosure to indicate that this PMF CRU will lose access to the Y fabric as well.
- [ ] On this worksheet, trace pathways of communication through all routers and ServerNet cables that connect to the CRU you plan to remove. Circle all components and enclosures that will be affected by that CRU removal.

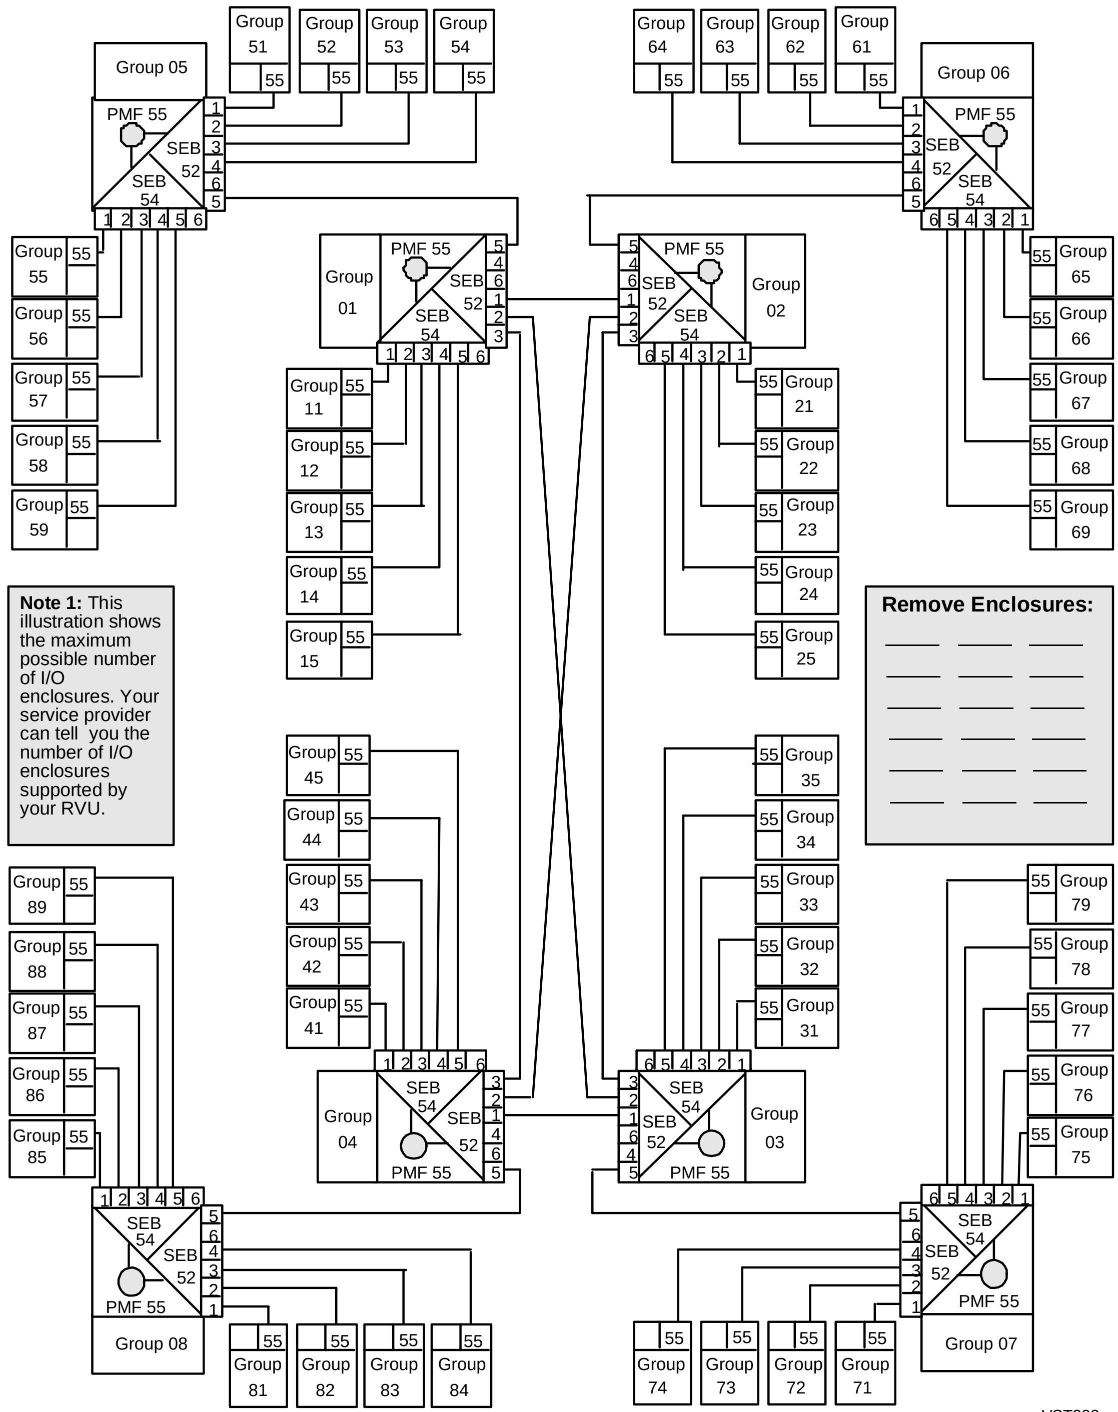

VST388.vsd

D

Stopping Devices and Processes

This appendix gives examples of identifying devices and processes before you move or stop them.

As stated in Step 5. Prepare and Stop Devices and Processes on page 4-8, you need to identify all devices and processes before you can stop or move them. This important step in system reduction must be done successfully if the donor system is to continue to function well after the enclosures have been physically removed.

This appendix contains some examples of how to carry out this step. Use these examples as a guide, along with the manuals for the remaining subsystems, to plan the specific commands needed for your system reduction. The procedures for identifying, preparing, and stopping these devices are subsystem-specific.

Based on Worksheet 8, the possible devices are listed in Table D-1 along with their corresponding manuals and a reference to an example if any.

Table D-1. References and Examples (page 1 of 2)

Device	Subsystem	Manual	Example on Page
Internal disk drive	Storage	*SCF Reference Manual for the Storage Subsystem*	D-2
External tape drive	Storage	*SCF Reference Manual for the Storage Subsystem* (if attached through PMF CRU or IOMF CRU)	--
		6760 ServerNet/DA Manual (if attached through ServerNet device adapter)	
Adapters (in general)	--	--	D-5
E4SA, GESA, and G4SA	Expand-over-IP	*Expand Configuration and Management Manual*	D-6
		LAN Configuration and Management Manual	
	IPX/SPX	*IPX/SPX Configuration and Management Manual*	--
	PAM	*PAM Configuration and Management Manual*	--
	SNAX	*SNAX/XF and SNAX/APN Configuration and Management Manual*	--

Table D-1. References and Examples (page 2 of 2)

Device	Subsystem	Manual	Example on Page
ServerNet/FX adapter	FOX	*ServerNet/FX Adapter Configuration and Management Manual*	--
6760 ServerNet device adapter	Storage	*6760 ServerNet/DA Manual*	--
ServerNet wide-area network concentrator (SWAN)	WAN	*WAN Subsystem Configuration and Management Manual*	--

Prepare and Stop Internal Disk Drives

In Step 4. Inventory Enclosures to Be Removed on page 4-7, you recorded the names of all disk drives in all the enclosures that you plan to remove from your system on one copy of Worksheet 8 for each enclosure. Examine those worksheets and determine which system-critical processes and applications use those disk drives so that you can move these processes and applications to enclosures that will remain in the donor system.

For example, if you are removing an enclosure that has a disk drive containing HP NonStop Kernel Open System Services (OSS) files or KMSF swap space, you need to move those OSS files or KMSF swap space to another enclosure that will not be removed. For more information, see the *Open System Services Management and Operations Guide* and the *Kernel-Managed Swap Facility (KMSF) Manual*, respectively.

Terminate or Move Processes to Other Internal Disk Drives

As part of the reduction procedure, several internal disk drives will be removed from the donor system. Before the reduction, terminate or move any dependent processes and files from those internal disk drives to disk drives that will remain in the system. Determining which processes and files are dependent is system-specific, so consult your system analyst and your service provider. Dependent processes can be of three types:

- Active processes currently using the disk drive
- Active processes not currently using the disk drive
- Processes not yet started that might need the disk drive

The following procedure identifies the active processes that are currently using an internal disk. This procedure might not generate a complete list of dependent processes.

Step	Action	Result
1.	At the TACL prompt, enter this command for each disk volume listed on Worksheet 8: `> FUP LISTOPENS` *`diskname`*`.*.*` For example, for the disk volume $OSS: `> FUP LISTOPENS $OSS.*.*`	You receive a list of processes that currently have files open on the specified disk.
2.	Record this information in the critical processes box on Worksheet 8.	
3.	For each process in the list, either terminate (that is, stop and delete) the process or copy the file it has open to a disk drive in an enclosure that will not be removed.	
4.	If you moved this file, reconfigure the process to open the copied file.	
5.	Repeat this procedure for each internal disk drive in the enclosure to be removed.	

Stop Internal Disk Drives

In each enclosure you remove, you stop all internal disk drives that are not already down.

1. At the SCF prompt, enter

   `-> STOP DISK (`*`diskname1,`*` [ , `*`disknamen`*` ]... )`

   You see a display similar to:

```
-> STOP DISK ($DATA01, $DATA02, $DATA03)
STORAGE W01007 The STOP DISK \TROLL.$DATA01 operation will cause
               \TROLL.$DATA01 to be INACCESSIBLE to user processes.
Do you want to STOP DISK \TROLL.$DATA01 (Y/[N])
```

2. Enter *Y*.

   If this disk contains open files, you receive a message similar to:

```
STORAGE W01008 DISK \TROLL.$DATA03 contains 2 open files.
Are you sure you want to STOP DISK $DATA03 (Y/[N])
```

   In this example, $DATA03 has two open files.

   - If you enter *Y*, you terminate access to the disk. Results in this case are system-dependent and might be unpredictable.
   - If you enter *N*, you can then analyze what process is accessing this disk and take appropriate action. Consult your system analyst and your service provider.

3. Verify that the disk drives have been stopped by entering this SCF command for each PMF CRU or IOMF CRU in each enclosure to be removed:

   -> STATUS ADAPTER $ZZSTO.*#device.group.module.slot,* DETAIL

   For example, for the IOMF CRU in slot 55 of I/O enclosure group 71:

   ```
   -> STATUS ADAPTER $ZZSTO.#IOMF.GRP-71.MOD-1.SLOT-55, DETAIL
   ```

   For the PMF CRU in slot 50 of processor enclosure group 07:

   ```
   -> STATUS ADAPTER $ZZSTO.#PMF.GRP-7.MOD-1.SLOT-50, DETAIL
   ```

   In this example, you see a display similar to:

```
STORAGE - Detailed Status ADAPTER \TROLL.$ZZSTO.#PMF.GRP-7.MOD-1.SLOT-50
 Location......... (7,1,50)        Number of SACs... 3
 Part ID.......... 127764          POST Result...... PASSED
 Power-1.......... ON              Power-2.......... ON
 Revision Level... A13-16          Status........... PRESENT
 Tracking Number.. D0C5D6          Vendor ID........

 ServerNet Addressable Controllers:

   SAC \TROLL.$ZZSTO.#PMF.SAC-1.GRP-7.MOD-1.SLOT-50:
    Subdevice...... 1              Node Class........ scsi
    Node ID........ 14             Node Type......... sc01
    Pack Class..... icpk           Pack Instance..... 1
    POST Result.... PASSED         SAC ID............ 6
    SAC Type....... SCSI           Side.............. X-Fabric
    Status......... PRESENT        Usage............. DISK
    Vendor ID...... 0

     Configured Devices ( group: 7 , module: 1  ):
       Slot  Name           State      Substate         Primary  Backup
                                                        PID      PID
          2  $DATA01-MB     STOPPED                      12,281    13,269
          4  $DATA02-P      STOPPED                      12,284    13,272
          6  $DATA03-MB     STARTED                      12,279    13,266

   SAC \TROLL.$ZZSTO.#PMF.SAC-2.GRP-7.MOD-1.SLOT-50:
    Subdevice...... 2              Node Class........ scsi
    Node ID........ 14             Node Type......... sc01
    Pack Class..... icpk           Pack Instance..... 1
    POST Result.... PASSED         SAC ID............ 6
    SAC Type....... SCSI           Side.............. X-Fabric
    Status......... PRESENT        Usage............. DISK
    Vendor ID...... 0

     Configured Devices ( group: 7 , module: 1  ):
       Slot  Name           State      Substate         Primary  Backup
                                                        PID      PID

         1  $DATA01-P   STOPPED                 12,281    13,269
          5  $DATA03-P     *STARTED                      12,279    13,266
```

   In this example display, the status of the disk drives appears in boldface. $DATA01 and $DATA03 are mirrored disks. $DATA02 is nonmirrored. $DATA03 is listed as

STARTED. Do not proceed with system reduction until all devices in the enclosure you are removing are listed as STOPPED.

Repeat this procedure to stop internal disk drives for each internal disk drive in each enclosure to be removed.

Delete Internal Disk Drives

Use the SCF DELETE DISK command to delete internal disks from the system configuration:

```
-> DELETE DISK ($DATA01, $DATA02, $DATA03)
```

From the OSM Service Connection or TSM Service Application tree pane, every internal disk you deleted should be listed as an Unknown Disk CRU.

Prepare and Stop Adapters

In Step 4. Inventory Enclosures to Be Removed on page 4-7, you recorded the names of all adapters in all the enclosures you are removing from your system on one copy of Worksheet 8 for each enclosure. Now examine those worksheets to determine how removal of an enclosure will affect the operations of your system.

In general, you are answering the following questions:

- What subsystem is used to control each adapter?
- What subsystem entities are subordinate to each adapter?
- What clients use those entities?

You will stop all adapters in each enclosure you will remove.

Procedures for stopping any device or subsystem are specific to that device or subsystem. For information about how to stop devices or subsystems, see the appropriate manual. In some cases, a process running in an enclosure that will not be removed might use hardware in an enclosure that will be removed.

In general, you will perform the following tasks within each subsystem:

1. Identify the subsystems that use adapters.
2. In each enclosure to be removed, stop as many I/O processes as you can.
3. Use subsystem-specific commands to stop the devices controlled by any I/O processes within that subsystem.

Prepare and Stop an E4SA

The next examples show how to stop an Ethernet 4 ServerNet adapter (E4SA). These examples:

- Use SCF unless otherwise noted
- Show part of the removal of group 02 from system \TROLL
- Show an E4SA that was configured to use only Expand-over-IP.
- Are only examples and might not be appropriate for your system

Note. Other subsystems besides Expand-Over-IP that an E4SA can use are SNAX over Ethernet, Port Access Method (PAM), and IPX/SPX.

This example screen shows all associations of Physical Interface (PIF) E4SA02A* to Logical Interface (LIF) LAN02* in the entire system. The LIF of the E4SA used in this example appears in boldface.

```
-> INFO LIF $ZZLAN.*

SLSA Info LIF

Name               PIF              MAC Address       Type
$ZZLAN.LAN020A  E4SA02A.0.A    08:00:8E:00:CF:5F  Ethernet
$ZZLAN.LAN020B  E4SA02A.0.B    08:00:8E:00:CF:60  Ethernet
$ZZLAN.LAN021A  E4SA02A.1.A    08:00:8E:00:CF:61  Ethernet
$ZZLAN.LAN021B  E4SA02A.1.B    08:00:8E:00:CF:62  Ethernet
$ZZLAN.LANX      MIOE0.0.A     08:00:8E:00:E3:77  Ethernet
$ZZLAN.LANY      MIOE1.0.A     08:00:8E:00:E3:79  Ethernet
```

This example screen shows in boldface the TCP/IP process that corresponds to the LIF LAN020A.

```
-> LISTOPENS LIF $ZZLAN.LAN020A

SLSA Status LIF \TROLL.$ZZLAN.LAN020A

Openers  PPID      BPID    Format     Filter
$ZTC20   ( 2,300) ( 3,296) Ethernet  $ZTC20.#SN1.BR0
$ZTC20   ( 2,300) ( 3,296) Ethernet  $ZTC20.#SN1.MC2
$ZTC20   ( 2,300) ( 3,296) Ethernet  $ZTC20.#SN1.IP0
$ZTC20   ( 2,300) ( 3,296) Ethernet  $ZTC20.#SN1.AR0
```

This example screen shows all open associations of the TCP/IP $ZTC20 process to other processes. For example, $ZTN20, the Telserv process, appears in boldface.

```
-> LISTOPENS PROCESS $ZTC20

TCPIP Listopens PROCESS \TROLL.$ZTC20

Openers              PPID      BPID     PLFN    BLFN    Protocol  Lport
$ZTN20               2,301              8       0       TCP       telnet
$ZPT20               0,325              4       0       TCP       echo
$ZTN20               2,301              9       0       TCP       telnet
$ZPT20               0,325              5       0       TCP       finger
$ZPT20               0,325              6       0       TCP       ftp
$ZTN20               2,301              7       0       TCP       telnet
$ZTN20               2,301              3       0       TCP       telnet
$ZPM20               2,302              3       0       UDP       111
$ZPM20               2,302              4       0       TCP       111
$ZBK20               1,41               3       0       TCP       795
$ZBK20               1,41               4       0       TCP       800
$ZTN20               2,301              10      0       TCP       telnet
$ZTN20               2,301              4       0       TCP       telnet
$ZNET                0,24               4       0       #ZSPI     *
$ZTN20               2,301              5       0       TCP       telnet
$ZTN20               2,301              6       0       TCP       telnet
```

This example screen shows all open Telserv sessions that the TCP/IP $ZTC20 process has:

```
-> LISTOPENS PROCESS $ZTN20

TELSERV Listopens PROCESS \TROLL.$ZTN20

Window    Openers            PCPU PPIN PLFN BCPU BPIN BLFN   Servic
#PTCZA3G  \TROLL   $Z02F     7    20   1    N/A  N/A  N/A    TACL
#PTCZA3V  \TROLL   $Z078     3    23   1    N/A  N/A  N/A    TACL
#PTCZA3Y  \TROLL   $Z09Q     6    27   1    N/A  N/A  N/A    TACL
#PTCZA3U  \TROLL   $Z06R     2    45   1    N/A  N/A  N/A    TACL
#PTCZA3R  \TROLL   $Z05F     7    25   1    N/A  N/A  N/A    TACL
#PTCZA3S  \TROLL   $Z06D     0    83   1    N/A  N/A  N/A    TACL
#PTCZA3Y  \TROLL             6    40   2    N/A  N/A  N/A    TACL
#PTCZA3W  \TROLL   $Z07F     4    26   1    N/A  N/A  N/A    TACL
#PTCZA3Y  \TROLL             6    40   3    N/A  N/A  N/A    TACL
#ZSPI     \TROLL   $ZNET     0    24   6    N/A  N/A  N/A    ZSPI
```

This example screen shows the TACL commands you can use to trace processes started from a TCP/IP session. The processes highlighted in bold type might be affected by the removal of this E4SA.

```
TACL> STATUS *, TERM $ZTN20.#PTCZA3G
Process    Pri PFR %WT   Userid   Program file        Hometerm
$Z02F 7,20 169     005  255,255  SYSTEM.SYS00.TACL  $ZTN20.#PTCZA3G

Process    Pri PFR %WT   Userid   Program file        Hometerm
$Z02F 6,45 168  P   000  255,255  SYSTEM.SYS00.FUP  $ZTN20.#PTCZA3Y
```

The previous five example screens show one session on one LIF on one E4SA.

On this LIF, a number of processes are affected when you stop the E4SA. Each of the following processes might need to be moved or stopped before you continue with the system reduction procedure:

- Processes that access the LIF through the following subsystems: Expand-Over-IP, PAM, SNAX over Ethernet, and IPX/SPX
- Direct openers of these same subsystems
- Clients of those direct openers

Determining this and similar information on any system is highly system-specific. Your procedures for determining critical processes in the enclosure you will remove might vary. Consult your system analyst.

Stop the E4SA Adapter

After you have identified the processes that need to be moved or stopped, enter the appropriate SCF commands to stop the E4SA adapter, its LIF, and the TCP/IP process. Depending on the subsystem, use the STOP command first, followed by the ABORT command if necessary.

In this example, the SCF STOP command performs an orderly shutdown of processes:

```
1-> STOP LIF $ZZLAN.LAN02*
2-> STOP ADAPTER $ZZLAN.E4SA02A, SUB ALL
3-> STOP PROCESS $ZTC20
```

If a process is still running when you try to stop it, SCF can return a message like:

```
SLSA E00012 Object $ZZLAN.LAN21 has registered client(s) and
cannot stop.
```

If all components have not been stopped and you have determined that they can be forcibly stopped, use the ABORT command:

```
1-> ABORT LIF $ZZLAN.LAN02*
2-> ABORT ADAPTER $ZZLAN.E4SA02A, SUB ALL
3-> ABORT PROCESS $ZTC20
```

Delete the E4SA Adapter

You can use these SCF commands to delete the #E4SA20A adapter along with its associated LIFs, PIFs, and SACs from the system configuration.

```
5-> DELETE LIF $ZZLAN.LAN020A
6-> DELETE LIF $ZZLAN.LAN020B
7-> DELETE LIF $ZZLAN.LAN021A
8-> DELETE LIF $ZZLAN.LAN021B
 9-> DELETE ADAPTER $ZZLAN.E4SA02A, SUB ALL
```

From the OSM Service Connection or TSM Service Application tree pane, the icon of every device that has been deleted should be yellow, and the Attributes tab should show the device as Not Configured.

Safety and Compliance

Regulatory Compliance Statements

The following warning and regulatory compliance statements apply to the products documented by this manual.

FCC Compliance

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.

Any changes or modifications not expressly approved by Hewlett Packard Computer Corporation could void the user’s authority to operate this equipment.

CISPR Compliance

This equipment complies with the requirements of CISPR 22 (EN 55 022) for Class A Information Technology Equipment (ITE). In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

Canadian Compliance

This class A digital apparatus meets all the requirements of the Canadian Interference-Causing Equipment Regulations.

Cet appareil numérique de la classe A respecte toutes les exigences du Règelment sur le matériel brouilleur du Canada.

Korea MIC Compliance

A급 기기 (업무용 정보통신기기)

이 기기는 업무용으로 전자파적합등록을 한 기기이오니 판매자 또는 사용자는 이 점을 주의하시기 바라며, 만약 잘못판매 또는 구입하였을 때에는 가정용으로 교환하시기 바랍니다.

Taiwan (BSMI) Compliance

警告使用者:

這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

Japan (VCCI) Compliance

この装置は、情報処理装置等電波障害自主規制協議会（ＶＣＣＩ）の基準に基づくクラスＡ情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

This is a Class A product based on the standard or the Voluntary Control Council for Interference by Information Technology Equipment (VCCI). If this equipment is used in a domestic environment, radio disturbance may occur, in which case the user may be required to take corrective actions.

DECLARATION OF CONFORMITY

Supplier Name:	**HP COMPUTER CORPORATION**	
Supplier Address:	**HP Computer Corporation,** **NonStop Enterprise Division** 10333 Vallco Parkway Cupertino, CA 95014 USA	***Represented in the EU By:*** **Hewlett Packard EMEA GmbH** P.O. Box 81 02 44 81902 Munich Germany

Declares under our sole responsibility that the following product

Product Name:	**NonStop S-series server**
Regulatory Model	**CPTOF-0301**
Product Model No:	**S88000, S86000, S76000, S7600, S74000, S7400, S78000, S7800**

Conforms to the following normative European and International Standards.

Product Safety:	**EN60950:2003**	**(IEC 60950-1 1st Edition)**
Electromagnetic Compatibility:	**EN 55022:1998**	**Radiated and Conducted Emission**
	EN 61000-3-2:2000	**Harmonic Current Emission**
	EN 61000-3-3 +A1:2001	**Voltage Fluctuation and Flicker**
	EN 55024:1998	**EMC Immunity**

Following the provisions of the normative European Council Directives:

EMC Directive 89/336/EEC (including amendments)

Low Voltage Directive 73/23/EEC (amended by 93/68/EEC)

Supplementary Information:

Safety: **Protection Class I, Pollution Degree II**
Emissions: **EMC Class A**
Year Assessed/First Production: 2004

Product conformance to cited product specifications is based on sample (type) testing, evaluation, or assessment at Hewlett Packard's compliance laboratories in Cupertino, California or at accredited laboratories accepted by European Union Notified and Competent Bodies.

Charles Denning
Manager, Hardware Product Assurance
NonStop Enterprise Division
Cupertino, California

Consumer Safety Statements

Customer Installation and Servicing of Equipment

The following statements pertain to safety issues regarding customer installation and servicing of equipment described in this manual.

- Keep door closed for normal operation.
- Batteries must be disposed of in compliance with local ordinances.

△ **Caution.** After unplugging the fan’s power cable from the PMCU, wait for the fan to stop spinning before performing the next servicing step.

△ **Caution.** Be careful not to drop the mounting screws into the fan unit.

▲ **WARNING.** Do not plug in the fan’s power cable until after you have mounted the fan. The fan blades begin to rotate as soon as the fan’s power cable is plugged into the PMCU.

▲ **WARNING.** Do not touch the DC power connector on the S7400/S7600/Sxx000 PMF CRU after you have unplugged the DC power cable from the CRU. It is possible to incur a severe energy hazard for as long as fifteen (15) seconds after unplugging the cable. You can make sure the CRU is safe to remove by testing (by using a volt meter) before touching the pins on the DC power connector.

Applicability of Procedures, Instructions, and Examples

▲ **WARNING.** The procedures and instructions and the examples of procedures and instructions found in this manual apply only to the system enclosures and components found in NonStop S-series servers which are described in the NonStop S-series server documentation. Do not attempt to use these instructions with any other NonStop system equipment. You can determine that your system is a NonStop S-series server by reading the label that appears inside each system enclosure on the CRU in slot 50 or 55. The label includes the product number 195x, 196x, 197x, or 198x. Read the label before proceeding with any procedure or instruction found in this manual.

Consignes de sécurité à l'intention du client

Installation et entretien du système par le client

Les consignes de sécurité qui suivent concernent l'installation et l'entretien par le client du système décrit dans le présent manuel.

- Garder la porte fermée pendant le fonctionnement normal du système.
- Jeter les piles usagées conformément au règlement local en vigueur.

△ **Attention.** Après avoir débranché le cordon d'alimentation du ventilateur de l'unité PMCU, attendre l'arrêt complet de l'hélice avant de passer à l'étape suivante.

△ **Attention.** Prendre soin de ne pas laisser tomber les vis de fixation dans le ventilateur.

▲ **MISE EN GARDE.** Attendre d'avoir remonté le ventilateur avant d'en rebrancher le cordon d'alimentation. En effet, l'hélice du ventilateur se met à tourner dès que l'on relie le ventilateur à l'unité PMCU.

▲ **MISE EN GARDE.** Ne pas toucher au connecteur d'alimentation C.C. de l'unité PMF remplaçable par le client du S7400/S7600/Sxx000 après avoir débranché le cordon d'alimentation C.C. de cette unité. Des risques de chocs électriques dangereux peuvent subsister pendant quinze (15) secondes après le débranchement du câble. Lors de la dépose de l'unité, on pourra confirmer l'absence de danger avec un voltmètre avant de toucher les broches du connecteur d'alimentation C.C.

Applicabilité des procédures, des directives et des exemples

▲ **MISE EN GARDE.** Les procédures et directives ainsi que les exemples y afférents contenus dans le présent manuel concernent uniquement les boîtiers et pièces des serveurs NonStop série S décrits dans la documentation connexe. Elles ne sont applicables à aucun autre système NonStop. Pour savoir si votre serveur est un NonStop de la série S, consultez l'étiquette figurant à l'intérieur du boîtier du système sur l'unité remplaçable par le client installée dans l'emplacement 50 ou 55. Cette étiquette indique le nombre de produit, 195x, 196x, 197x, ou 198x. Consultez cette étiquette avant d'exécuter l'une quelconque des procédures ou directives prescrites dans le présent manuel.

Verbraucher-Sicherheitsangaben

Geräteinstallation und -wartung durch den Kunden

Die folgenden Angaben betreffen Sicherheitsfragen in Hinsicht auf die Geräteinstallation und -wartung durch den Kunden, wie sie in diesem Handbuch beschrieben werden.

- Tür für normalen Betrieb geschlossen lassen.
- Batterien müssen in Übereinstimmung mit örtlichen Vorschriften beseitigt werden.

△ **Vorsicht.** Nach dem Trennen des Gebläse-Stromversorgungskabels von der PMCU, vor dem Durchführen des nächsten Wartungsschritts auf das Ende der Drehbewegung des Gebläses warten.

△ **Vorsicht.** Achten Sie bitte darauf, daß die Befestigungsschrauben nicht in die Gebläseeinheit fallen.

▲ **WARNUNG.** Stecken Sie das Gebläse-Stromversorgungskabel erst nach dem Befestigen des Gebläses ein. Die Gebläseflügel beginnen sich zu drehen, sobald das Gebläse-Stromversorgungskabel in die PMCU eingesteckt wird.

▲ **WARNUNG.** Den Gleichstrom-Steckverbinder an der vom Kunden austauschbaren PMF-Einheit des S7400/S7600/Sxx000 nicht berühren, nachdem das Gleichstromkabel von der vom Kunden austauschbaren Einheit abgezogen wurde. Bis zu fünfzehn (15) Sekunden nach Abziehen des Kabels besteht starke Elektroschockgefahr. Durch Prüfen (mit einem Spannungsmesser) läßt sich vor Berührung der Stifte am Gleichstrom-Steckverbinder feststellen, ob die vom Kunden austauschbare Einheit gefahrlos entfernt werden kann.

Anwendbarkeit von Verfahren, Anleitungen und Beispielen

▲ **WARNUNG.** Die in diesem Handbuch enthaltenen Verfahren und Anleitungen und die Beispiele von Verfahren und Anleitungen gelten nur für die in den NonStop-S-Serie-Servern vorgefundenen Systemgehäuse und Komponenten, die in der Dokumentation der NonStop-S-Serie-Server beschrieben werden. Versuchen Sie nicht, diese Anleitungen mit anderen NonStop-System-Geräten zu benutzen. Sie können feststellen, daß Ihr System ein NonStop-S-Serie-Server ist, indem Sie den Aufkleber lesen, der sich im Innern eines jeden Systemgehäuses befindet in der vom Kunden austauschbaren Einheit in Steckplatz 50 oder 55. Der Aufkleber soll die Produktnummer 195x, 196x, 197x, oder 198x angeben. Lesen Sie den Aufkleber, bevor Sie mit irgendeiner der in diesem Handbuch enthaltenen Verfahren oder Anleitungen fortfahren.

Declaraciones sobre la seguridad del consumidor

Instalación y servicio al equipo por el consumidor

Las siguientes declaraciones tienen que ver con aspectos de seguridad relacionados con la instalación y servicio al equipo por el consumidor, y que se describen en este manual.

- Mantenga la puerta cerrada durante la operación normal del equipo.
- Las baterías (pilas) deben desecharse cumpliendo con los reglamentos locales.

△ **Precaución.** Después de desconectar el cable de alimentación del ventilador del PMCU, espere a que el ventilador deje de girar antes de realizar el siguiente paso de servicio.

△ **Precaución.** Tenga cuidado y no deje caer los tornillos de montaje dentro de la unidad del ventilador.

▲ **ADVERTENCIA.** No enchufe el cable de alimentación del ventilador sino hasta después de haber montado el ventilador. Las aspas del ventilador comienzan a girar tan pronto como se enchufe el cable de alimentación del ventilador en el PMCU.

▲ **ADVERTENCIA.** No toque el conector de alimentación de corriente directa en la unidad reemplazable por el cliente S7400/S7600/Sxx000 PMF después de haber desenchufado el cable de alimentación de corriente directa de dicha unidad. Es posible incurrir en un peligro grave debido a la energía eléctrica hasta por unos quince (15) segundos después de haber desenchufado el cable. Usted puede asegurarse de que la unidad reemplazable por el cliente ya se pueda quitar sin problema alguno probándola (con un voltímetro) antes de tocar las clavijas del conector de alimentación de corriente directa.

Forbrugersikkerhedsmeddelelser

Installation og service af udstyr der udføres af kunden

De følgende meddelelser vedrører sikkerheden angående installation og service af udstyr, der udføres af kunden, som beskrives i denne brugerhåndbog.

- Hold lugen lukket under normal drift.
- Batterierne skal kasseres i overensstemmelse med lokale vedtægter.

△ **Forsigtig!** Når du har afbrudt ventilatorens netkabel fra PMCU'en, skal du vente til ventilatoren standser helt, før du udfører det næste servicetrin.

△ **Forsigtig!** Pas på, at du ikke taber monteringsskruerne ned i ventilatoren.

▲ **ADVARSEL!** Tilslut ikke ventilatorens netkabel indtil efter, at du har monteret ventilatoren. Ventilatorvingerne begynder at dreje, så snart ventilatorens netkabel tilsluttes PMCU'en.

▲ **ADVARSEL!** Berør ikke jævnstrøms forbindelsesstikket på S7400/S7600/Sxx000 PMF-enheden, der kan udskiftes af kunden, når du har afbrudt jævnstrøms netkablet fra enheden, der kan udskiftes af kunden. Der er risiko for alvorligt stød i op til 15 sekunder efter, at kablet afbrydes. Du kan sikre dig, at det er sikkert at fjerne enheden, der kan udskiftes af kunden, ved at teste den (med et voltmeter), før du berører stifterne på jævnstrøms forbindelsesstikket.

Veiligheidsinstructies voor de consument

Installatie en onderhoud van apparatuur door de klant

De volgende veiligheidsinstructies betreffen de installatie en het onderhoud door de klant van de in deze handleiding beschreven apparatuur.

- Houd bij normaal bedrijf de deur gesloten.
- Batterijen moeten overeenkomstig de plaatselijke voorschriften worden weggegooid.

△ **Opgelet.** Nadat de stroomkabel van de ventilator van de PMCU is losgekoppeld, moet u wachten tot de ventilator stilstaat voordat u de volgende onderhoudsstap uitvoert.

△ **Opgelet.** Pas op dat u de montageschroeven niet in de ventilator laat vallen.

▲ **WAARSCHUWING.** Sluit de stroomkabel van de ventilator pas aan nadat u de ventilator gemonteerd hebt. De ventilatorbladen beginnen te draaien zodra de stroomkabel van de ventilator op de PMCU wordt aangesloten.

▲ **WAARSCHUWING.** Raak de gelijkstroomconnector op het S7400/S7600/Sxx000 PMF apparaat dat door de klant kan worden vervangen, niet aan nadat de gelijkstroomkabel is losgekoppeld van het door de klant vervangbare apparaat. Een ernstig energierisico kan nog wel vijftien (15) seconden nadat de kabel is losgekoppeld, aanwezig zijn. Controleer of het door de klant vervangbare apparaat veilig kan worden verwijderd door (met een spanningsmeter) een test uit te voeren voordat u de pennen op de gelijkstroomconnector aanraakt.

Käyttöturvaa koskevia huomautuksia

Asiakkaan suorittama laiteasennus ja huolto

Seuraavat huomautukset koskevat turvallisuusnäkökohtia, jotka asiakkaan täytyy ottaa huomioon tässä käsikirjassa kuvattuja laiteasennuksia ja huoltotoimenpiteitä suoritettaessa.

- Kansi täytyy pitää suljettuna normaalin käytön aikana.
- Paristot täytyy hävittää paikallisten säädösten mukaisesti.

△ **Varoitus.** Kytkettyäsi tuulettimen virtajohdon irti PMCU:sta odota, että tuuletin lakkaa pyörimästä, ennen kuin suoritat seuraavan huoltovaiheen.

△ **Varoitus.** Varo pudottamasta kiinnitysruuveja tuulettimeen.

▲ **VAARA.** Älä yhdistä tuulettimen virtajohtoa ennen kuin olet asentanut tuulettimen paikoilleen. Tuulettimen siivet alkavat pyöriä heti, kun tuulettimen virtajohto on yhdistetty PMCU:hun.

▲ **VAARA.** Älä kosketa S7400/S7600/Sxx000 PMF vaihto-osan virtaliitintä heti sen jälkeen kun olet irrottanut vaihto-osan tasavirtakaapelin. Vaarallista jännitettä voi olla jäljellä vielä viidentoista (15) sekunnin kuluttua kaapelin irrottamisesta. Jännitemittarilla on mahdollista testata, onko vaihto-osan tasavirtaliittimen pistokkeiden koskettaminen turvallista.

Veiligheidsinstructies voor de consument

Installatie en onderhoud van apparatuur door de klant

De volgende veiligheidsinstructies betreffen de installatie en het onderhoud door de klant van de in deze handleiding beschreven apparatuur.

- Houd bij normaal bedrijf de deur gesloten.
- Batterijen moeten overeenkomstig de plaatselijke voorschriften worden weggegooid.

△ **Opgelet.** Nadat de stroomkabel van de ventilator van de PMCU is losgekoppeld, moet u wachten tot de ventilator stilstaat voordat u de volgende onderhoudsstap uitvoert.

△ **Opgelet.** Pas op dat u de montageschroeven niet in de ventilator laat vallen.

▲ **WAARSCHUWING.** Sluit de stroomkabel van de ventilator pas aan nadat u de ventilator gemonteerd hebt. De ventilatorbladen beginnen te draaien zodra de stroomkabel van de ventilator op de PMCU wordt aangesloten.

▲ **WAARSCHUWING.** Raak de gelijkstroomconnector op het S7400/S7600/Sxx000 PMF apparaat dat door de klant kan worden vervangen, niet aan nadat de gelijkstroomkabel is losgekoppeld van het door de klant vervangbare apparaat. Een ernstig energierisico kan nog wel vijftien (15) seconden nadat de kabel is losgekoppeld, aanwezig zijn. Controleer of het door de klant vervangbare apparaat veilig kan worden verwijderd door (met een spanningsmeter) een test uit te voeren voordat u de pennen op de gelijkstroomconnector aanraakt.

Misure precauzionali per i clienti

Installazione e manutenzione del sistema da parte del cliente

Le seguenti misure precauzionali riguardano l'installazione e la manutenzione da parte del cliente del sistema descritto nel presente manuale.

- Mantenere la porta chiusa durante il funzionamento normale del sistema.
- Lo smaltimento delle batterie usate deve essere effettuato secondo la normativa locale.

△ **Avvertenza.** Dopo aver scollegato il cavo di alimentazione del ventilatore dall'unità PMCU, attendere che il ventilatore smetta di girare prima di procedere con la manutenzione.

△ **Avvertenza.** Fare attenzione a non lasciar cadere le viti di montaggio nell'unità del ventilatore.

▲ **ATTENZIONE.** Inserire il cavo di alimentazione del ventilatore solamente dopo aver montato il ventilatore stesso. Le pale del ventilatore iniziano a ruotare non appena il cavo di alimentazione del ventilatore viene inserito nell'unità PMCU.

▲ **ATTENZIONE.** Non toccare il connettore di corrente continua sull'unità sostituibile dal cliente S7400/S7600/Sxx000 PMF dopo aver disinserito il cavo di corrente continua dall'unità stessa. Un serio rischio elettrico può perdurare fino a 15 secondi dopo aver disinserito il cavo. È possibile controllare con un voltmetro che l'unità sostituibile dal cliente possa essere rimossa con sicurezza prima di toccare i piedini del connettore di corrente continua.

Informações de segurança para os consumidores

Instalação e manutenção do equipamento pelo cliente

As seguintes informações se referem a questões de segurança relacionadas à instalação e manutenção, pelo cliente, do equipamento descrito neste manual.

- Para garantir o funcionamento normal, mantenha a porta fechada.
- As pilhas usadas devem ser descartadas de acordo com as leis locais.

△ **Cuidado.** Após desligar o cabo de alimentação do ventilador do PMCU, espere o ventilador parar antes de prosseguir para a etapa seguinte da manutenção.

△ **Cuidado.** Tenha cuidado para não deixar que os parafusos de montagem caiam dentro do ventilador.

▲ **AVISO.** Só ligue o cabo de alimentação do ventilador após ter montado o ventilador. As lâminas do ventilador começam a girar logo que o cabo de alimentação do ventilador é ligado ao PMCU.

▲ **AVISO.** Não toque no conector de alimentação DC da Unidade Substituível pelo Cliente (USC) S7400/S7600/Sxx000 PMF logo após ter desligado o cabo de alimentação DC da USC. Existe a possibilidade de ocorrer um acidente elétrico sério até quinze (15) segundos após o desligamento do cabo. Assegure-se de que é seguro remover a USC testando-a (ou usando um medidor de voltagem), antes de tocar nos pinos do conector de alimentação DC.

Informações de segurança para os consumidores

Instalação e manutenção do equipamento pelo cliente

As seguintes informações referem-se a questões de segurança relacionadas à instalação e manutenção, pelo cliente, do equipamento descrito neste manual.

- Para garantir o funcionamento normal, mantenha a porta fechada.
- As pilhas usadas devem ser descartadas de acordo com as leis locais.

△ **Cuidado.** Após desligar o cabo de alimentação do ventilador do PMCU, espere que o ventilador páre antes de prosseguir para a etapa seguinte da manutenção.

△ **Cuidado.** Tenha cuidado para não deixar que os parafusos de montagem caiam dentro do ventilador.

▲ **AVISO.** Ligue o cabo de alimentação do ventilador só depois ter montado o mesmo. As lâminas do ventilador começam a girar logo que o cabo de alimentação do ventilador é ligado ao PMCU.

▲ **AVISO.** Não toque no conector de alimentação DC da Unidade Substituível pelo Cliente (USC) S7400/S7600/Sxx000 PMF logo após ter desligado o cabo de alimentação DC da USC. Existe a possibilidade de ocorrer um acidente eléctrico sério até quinze (15) segundos após o desligamento do cabo. Assegure-se de que é seguro remover a USC testando-a (ou usando um medidor de voltagem), antes de tocar nos pinos do conector de alimentação DC.

Meddelanden beträffande konsumentsäkerhet

Kundutförd installation och service

De följande meddelandena beskriver säkerhetsföreskrifter för kundutförd installation och service av utrustning som beskrivs i denna manual:

- Dörren skall vara stängd under normal drift.
- Batterier måste kasseras i enlighet med lokala förordningar.

△ **Observera!** När du kopplat ur nätsladden till fläkten från PMCU-enheten måste du vänta tills fläktbladen har stannat innan du utför nästa servicemoment .

△ **Observera!** Se till att du inte tappar monteringsskruvarna in i fläkten.

▲ **VARNING!** Fläktbladen börjar snurra så snart fläktens nätsladd ansluts till PMCU-enheten. Du ansluter nätsladden till fläkten efter det att du har monterat fläkten.

▲ **VARNING!** Likströmskontakten som är placerad på S7400/S7600/Sxx000 PMF-enheten är strömförande upp till femton (15) sekunder efter att elkablen till enheten har blivit urkopplad. Du skyddar dig från strömstötar genom att testa om likströmskontakten fortfarande är strömförande (med hjälp av en voltmätare).

機器のカスタマー・インストレーションおよび保守

次の記述は、このマニュアルに述べられた機器のカスタマー・インストレーションおよび保守に関する安全性の問題に適合するものです。

通常のオペレーションではドアを閉める。

バッテリーは定められた法規に適合するものであること。

注：PMCU からファンの電源ケーブルを抜いた後は、ファンの回転が止まるまで待ってから、次の保守ステップに進むこと。

注：ファン・ユニットの中に装填スクリューを落とさないように注意する。

警告：ファンを装填するまでは、ファンの電源ケーブルを差し込まないこと。ファンの羽根は、PMCU にファンの電源ケーブルを差し込むと、直ちに回転を始めるため。

警告：CRU から DC 電源ケーブルを抜いた後で、S7400/S7600/Sxx000 の PMF CRU 上の DC 電源コネクターに触らないこと。ケーブルを抜いてから 15 秒間は、大きなエネルギー事故を引き起こす危険性があるため。CRU を安全に取り外すには、DC 電源コネクターのピンに触る前に（ボルト・メーターを使って）テストする。

用户安全使用说明

用户安装及使用

下列内容是手册中有关用户安装使用时的安全注意事项：

- 正常运行时请将门保持关闭。
- 设备电源线应与插座靠近，客户应容易接近插座。

△ **小心。** 从电源监控装置 (PMCU) 上拔下风扇电源插头进行下一步骤之前，应等待风扇完全停止转动。

△ **小心。** 不要把安装螺钉掉进风扇。

▲ **警惕。** 在风扇安装好之前，不要接通电源。电扇电源线一插上电源监控装置 (PMCU)，电扇叶片马上开始转动。

▲ **警惕。** 从用户替换电源 (CRU) 上拔下直流电源线之後，不要在 S7400/ S7600/ Sxx000 PMF 用户替换电源 (CRU) 上触摸直流电源接线板。在拔下电源後的十五秒内，仍有可能造成严重的电力事故。在触摸直流电源接线板上的插头之前，请用电压表确认用户替换电源 (CRU) 已安全切断。

Οδηγίες ασφαλείας του καταναλωτή

Εγκατάσταση και συντήρηση του εξοπλισμού από τον πελάτη

Οι παρακάτω οδηγίες αφορούν την ασφάλεια του πελάτη σχετικά με την εγκατάσταση και συντήρηση του εξοπλισμού του από τον ίδιο όπως αναφέρεται στα εγχειρίδια.

- **Η πόρτα πρέπει να είναι κλειστή για την κανονική λειτουργία του.**
- **Οι μπαταρίες πρέπει να αχρηστεύονται σύμφωνα με τους τοπικούς κανονισμούς.**

Προσοχή. Αφότου αποσυνδέσετε το καλώδιο ρεύματος του ανεμιστήρα από τη Μονάδα Ελέγχου και Ενδεικτικό Τροφοδοσίας (PMCU), περιμένετε να σταματήσει να γυρίζει ο ανεμιστήρας πριν συνεχίσετε την συντήρηση.

Προσοχή. Προσέχετε να μην σας πέσουν οι βίδες στήριξης μέσα στη μονάδα του ανεμιστήρα.

ΠΡΟΣΟΧΗ. Μην τροφοδοτείτε με ρεύμα τον ανεμιστήρα πριν τον εγκαταστήσετε. Τα πτερύγια του θα αρχίσουν να περιστρέφονται μόλις συνδέσετε το καλώδιο τροφοδοσίας του ανεμιστήρα στη Μονάδα Ελέγχου και Ενδεικτικό Τροφοδοσίας (PMCU).

S7400/
S7600/
Sxx000

ΠΡΟΣΟΧΗ. Μην αγγίζετε το βύσμα συνεχούς ρεύματος του PMF (Επεξεργαστή Πολλαπλής Λειτουργίας) της Αντικαταστάσιμης από τον Πελάτη Μονάδας (CRU) αφότου έχετε αποσυνδέσει το καλώδιο τροφοδοσίας από την Αντικαταστάσιμη από τον Πελάτη Μονάδα (CRU). Υπάρχει σοβαρός κίνδυνος ηλεκτροπληξίας μέχρι και δεκαπέντε (15) δευτερόλεπτα αφότου αποσυνδέσετε το καλώδιο τροφοδοσίας. Μπορείτε να σιγουρευτείτε για την ασφαλή αφαίρεση της Αντικαταστάσιμης από τον Πελάτη Μονάδας (CRU) δοκιμάζοντας την (με βολτόμετρο) πριν αγγίξετε τις ακίδες του βύσματος.

Glossary

3-phase. Describes a single power source with three output phases (A, B, and C). The phase difference between any two of the three phases or currents is 120 degrees.

3860 ATM 3 ServerNet adapter (ATM3SA). See ATM 3 ServerNet adapter (ATM3SA).

3861 Ethernet 4 ServerNet adapter (E4SA). See Ethernet 4 ServerNet adapter (E4SA).

3862 Token-Ring ServerNet adapter (TRSA). See Token-Ring ServerNet adapter (TRSA).

3863 Fast Ethernet ServerNet adapter (FESA). See Fast Ethernet ServerNet adapter (FESA).

3865 Gigabit Ethernet ServerNet adapter (GESA). See Gigabit Ethernet ServerNet adapter (GESA).

4619 disk drive. An 18-gigabyte, 15,000-rpm, small computer system interface (SCSI) disk drive for HP NonStop™ S-series servers running G06.06 and later release version updates (RVUs). The 4619 disk drive can coexist and operate with lower-capacity or lower-speed drives in the same storage subsystem module.

4637 disk drive. A 36-gigabyte, 10,000-rpm, small computer system interface (SCSI) disk drive for HP NonStop™ S-series servers running G06.06 and later release version updates (RVUs).

6740 ServerNet/FX adapter. See ServerNet/FX adapter.

6742 ServerNet/FX 2 adapter. See ServerNet/FX 2 adapter.

6760 ServerNet device adapter (ServerNet/DA). See ServerNet device adapter (ServerNet/DA).

6761 F-PIC. See fiber-optic plug-in card (F-PIC).

6762 S-PIC. See SCSI plug-in card (S-PIC).

6763 Common Communication ServerNet adapter (CCSA). See Common Communication ServerNet adapter (CCSA).

A. See ampere (A).

A0CINFO file. A distribution subvolume (DSV) file that contains information about a product and each of its files, including product and file dependencies, how the files are used and where they are placed, and which type of processor the product runs on. Every product and software product revision (SPR) to be managed by Distributed Systems Management/Software Configuration Manager (DSM/SCM) is distributed in a subvolume, and that subvolume must contain the product's A0CINFO file.

absolute pathname. An Open System Services (OSS) pathname that begins with a slash (/) character and is resolved beginning with the root directory. Contrast with relative pathname.

AC. See alternating current (AC).

accelerated mode. The operational environment in which Accelerator-generated RISC instructions execute. See also TNS mode and TNS/R native mode.

accelerated object code. The RISC instructions that result from processing a TNS object file with the Accelerator program.

accelerated object file. The object file that results from processing a TNS object file with the Accelerator program. An accelerated object file contains the original TNS object code, the accelerated object code and related address map tables, and any binder and symbol information from the original TNS object file.

Accelerator program. A program that processes a TNS object file and produces an accelerated object file. Most TNS object code that has been accelerated runs faster on TNS/R processors than TNS object code that has not been accelerated. The Accelerator program (AXCEL) is run prior to running the accelerated linker, XLLINK.

access mode. The form of file access permitted for a user or process.

ACL. See automatic cartridge loader (ACL).

ACS. See automated cartridge subsystem (ACS).

action. An operation that can be performed on a selected resource.

activation. The operator action of putting software into use after the software has been applied from the activation package to the target system.

activation package. A set of files containing product files, operator instructions, and instructions for applying the software on the target system. It consists of a header file containing the activation instructions and file attributes, multiple data files, Distributed Systems Management/Software Configuration Manager (DSM/SCM) control information, and Event Management Service (EMS) events.

AC transfer switch. A component of an HP NonStop™ Cluster Switch (model 6770) that provides access to dual AC power sources and the ability to switch between the two sources if one fails. The AC transfer switch draws power from its primary power source as long as it is available. If the primary source fails, the AC transfer switch is switched to draw power from the secondary power source.

adapter. See ServerNet adapter.

adapter cable. (1) A cable that connects components that have incompatible electrical interfaces. (2) For the ServerNet wide area network (SWAN) concentrator, one of four types of cable that can connect any of the six 50-pin WAN ports to one of the supported electrical interfaces (RS-232, RS-449, X.21, or V.35).

ADAPTER object type. The Subsystem Control Facility (SCF) object type for all adapters attached to your system.

address space. The memory locations to which a process has access.

ADE. See application development environment (ADE).

adjacent SP. A service processor (SP) that is directly connected through the ServerNet fabrics to the enclosure of a specified SP.

administrator. (1) For an HP NonStop™ system, the person responsible for the installation and configuration of a software subsystem on a NonStop node. Contrast with operator. (2) For an IBM system, the person responsible for the day-to-day monitoring and maintenance tasks associated with a software subsystem on an IBM node. (3) For a UNIX system, the owner of `/dev/console`. The administrator is responsible for the installation and configuration of all hardware and software within a node.

ADP. See Automated Data Processing (ADP).

ALLPROCESSORS paragraph. A required paragraph in the CONFTEXT configuration file that contains attributes defining the HP NonStop™ Kernel operating system image for all system processors. The ALLPROCESSORS paragraph follows the optional DEFINES paragraph.

alternate path. A path not enabled as the preferred path. An alternate path can become a primary path when a primary path is disabled.

alternating current (AC). An electric current having a waveform that regularly reverses in positive and negative directions. North American electrical power alternates 60 times/second (60 hertz). Contrast with direct current (DC).

amperage. Current-carrying capacity, expressed in amperes.

ampere (A). The unit of electrical current or rate of flow of electrons. One volt across one ohm of resistance causes a current flow of one ampere. A flow of one coulomb/second equals one ampere.

ANSI. The American National Standards Institute.

APE. See Application Program Examiner (APE).

API. See application program interface (API).

appearance side. The side of a system enclosure that contains, behind a door, disk customer-replaceable units (CRUs) and power monitor and control unit (PMCU) CRUs.

The appearance side is opposite the service side. System enclosures are typically arranged so that the appearance side is the most visible side. See also service side.

application binary interface (ABI). The conventions used to call functions and access global or external data.

application development environment (ADE). A set of methods and tools that are used throughout the lifecycle of an application project to design, code, and manage that project.

Application Program Examiner (APE). A tool used to browse through TNS object files that have been accelerated by the Accelerator. APE displays MIPS RISC code in addition to TNS code.

application program interface (API). A set of services (such as programming language functions or procedures) that are called by an application program to communicate with other software components. For example, an application program in the form of a client might use an API to communicate with a server program.

application-specific integrated circuit (ASIC). A custom-built integrated circuit (IC) used to perform highly specialized functions.

Apply. The Distributed Systems Management/Software Configuration Manager (DSM/SCM) action of executing the instructions contained in an activation package, such as placing new software on the target system and taking a snapshot of the new target system.

appropriate privileges. In the Open System Services (OSS) environment, an implementation-defined means of associating privileges with a process for function calls or function call options that need special privileges.

Archive. A set of unstructured files used to collect the software received onto the host system. Files received as input are placed in the Archive, and attributes of the files are stored in the host database. The planner specifies the Archive location in the Configuration Manager profile, using the Archive and Database Maintenance Interface.

Archive and Database Maintenance Interface. A block-mode interface run by a database or system administrator at both the host system and target systems to perform Distributed Systems Management/Software Configuration Manager (DSM/SCM) maintenance functions.

ASCII. American Standard Code for Information Interchange. A single-byte code set that uses only 7 of the 8 bits in a byte to represent each character. The ASCII code set contains the uppercase and lowercase characters of the U.S. English alphabet, some punctuation symbols, the digits 0 through 9, and some symbols and control characters. Because of its limited characters, and because the 8th bit is sometimes used in ASCII programs as a utility bit, the ASCII code set is not appropriate for use in international software.

ASIC. See application-specific integrated circuit (ASIC).

ASSIGN. An HP Tandem Advanced Command Language (TACL) command you can use to associate a file name with a logical file of a program or to assign a physical device to logical entities that an application uses.

assign message. Within Subsystem Control Facility (SCF), a message created by SCF for each ASSIGN command. A new process must request its assign message following receipt of the startup message. All assign messages set by the SCF ASSIGN command, plus the ones read from the HP Tandem Advanced Command Language (TACL) command interpreter, are passed to the new process.

assumed object. The object type or object name specified by a Subsystem Control Facility (SCF) ASSUME command. If an ASSUME command has been used to establish a default object type and fully qualified default object name, and if that object type and object name together refer to a valid object, `object-spec` can be omitted entirely from an SCF command, and the command is applied to the object known as the assumed object.

asynchronous wide area network (AWAN) servers. A local area network (LAN)-based communications device that provides (1) asynchronous connections to terminals, printers, and terminal emulators for HP NonStop S-series and K-series servers; (2) remote-access disk operating system (DOS), Windows, and Macintosh platforms; (3) VT-to-6530 protocol conversion; and (4) dial-out connections for LAN-attached DOS, Windows, and Macintosh platforms.

ATM3SA. See ATM 3 ServerNet adapter (ATM3SA).

ATM 3 ServerNet adapter (ATM3SA). A ServerNet adapter that provides access to Asynchronous Transfer Mode (ATM) networks from an HP NonStop™ S-series server. The 3860 ATM3SA supports the ATM User-Network Interface (UNI) specification over a 155-megabit/second (Mbps) OC-3 Sonet (Synchronous Optical Network) connection.

atomic. Behaving as a single, indivisible operation. For example, an atomic write operation on a file cannot write data that is interleaved with data from another, concurrent write operation on that file.

attachment. A file that contains information that augments the information in an incident report.

attribute. (1) For the Subsystem Control Facility (SCF), a characteristic of an entity. For example, two attributes of a process might be its program file and its user ID. An attribute is sometimes called a modifier. (2) In OSM and TSM client interfaces, a data item associated with a system or cluster resource. All attributes can be viewed, and some can be modified.

audit. A Distributed Systems Management/Software Configuration Manager (DSM/SCM) activity initiated by the operator at a target system that updates the target database with the fingerprints of all the files in a selected set of target subvolumes (TSVs).

authentication attributes. Security attributes of a process that do not change unless a successful reauthentication occurs or the super ID changes them. For Open System Services (OSS) processes, the authentication attributes include the login name, real user ID, real group ID, authentication system (node name), and group list.

authorization attributes. Security attributes of a process that can change through use of functions such as `setuid()` (or of Guardian procedures such as PROCESS_CREATE_) without reauthentication. For Open System Services (OSS) processes, the authorization attributes include the effective user ID, saved-set user ID, saved-set group ID, user audit flags, and effective user name.

authorization key. A password required for logging on to a modem. If you plan to allow dial-outs to a service provider, you must specify the authorization key of the service provider's modem during configuration of the OSM or TSM Notification Director.

automated cartridge subsystem (ACS). A type of tape library. Also known as automated cartridge system.

Automated Data Processing (ADP). The term used in the FIPS PUB 94 document to refer to computerized data processing equipment that is installed inside a computer room.

automatic cartridge loader (ACL). A device that stores multiple cartridge tapes and loads them automatically, one at a time, into a tape drive.

automatic configuration. The automatic assignment of magnetic disk attributes to an internal disk drive when it is inserted into a slot. Also known as *plug and play*.

averaging. A measurement method for determining the average value of alternating voltage and current waveforms. The averaging method involves sampling a waveform and averaging the samples over the period of one cycle.

AWAN. See asynchronous wide area network (AWAN) servers.

AXCEL. The command used to invoke the Accelerator on a TNS/R system.

back-end board (BEB). A circuit board that translates fiber-optic signals from a 3216 controller or 6760 ServerNet device adapter into small computer system interface (SCSI) commands and information for a tape drive. The BEB is housed in a cage-like sheet-metal enclosure and plugs into one of the 50-pin SCSI ports on the back of a tape drive customer-replaceable unit (CRU).

background process. In the Open System Services (OSS) environment, a process that belongs to a background process group.

background process group. In the Open System Services (OSS) environment, a process group that is both:

- Not a foreground process group
- A member of a session that has a connection with a controlling terminal

backout. The Distributed Systems Management/Software Configuration Manager (DSM/SCM) action of making the last configuration applied to the target system inaccessible and replacing it with the previous configuration.

backplane. A board that has connectors, on one or both sides of the board, into which circuit board assemblies plug. Backplanes are located behind card cages.

BACKUP. A utility for the HP NonStop™ servers that creates a backup copy of one or more disk files on magnetic tape. See also RESTORE.

backup processor. A processor running the HP NonStop™ Kernel operating system that communicates with the primary processor, allowing the processors to remain independent. A component failure in one processor has no effect on any other processor.

base computing platform. The minimum software implementation that is the foundation for the X/Open common applications environment (CAE).

base enclosure. An enclosure that is placed on the floor and can have other enclosures stacked on top of it. A base enclosure is installed on a frame base. Contrast with stackable enclosure.

base profile. In an X/Open compliant system, a minimum set of software components required to create a common applications environment.

battery load. The electrical current drain imposed on a battery.

BEB. See back-end board (BEB).

BIC. Backplane interconnect card. Not applicable to HP NonStop™ S-series servers. See ServerNet adapter.

BIND. A program invoked during system generation that creates TNS object (file code 100) system code files and system library files.

Binder. A programming utility that combines one or more compilation units' TNS object code files to create an executable TNS object code file for a TNS program or library. Used only with TNS object files.

Binder region. The region of a TNS object file that contains header tables for use by the Binder program.

binding. The operation of collecting, connecting, and relocating code and data blocks from one or more separately compiled TNS object files to produce a target object file.

bit-synchronous. A type of Open Systems Interconnection (OSI) Layer-2 protocol that uses synchronous transmission but does not require a character code to define terminal and line control sequences.

block. A grouping of one or more system enclosures that an HP NonStop™ S-series system recognizes and supports as one unit. A block can consist of either one processor enclosure, one I/O enclosure, or one processor enclosure with one or more I/O enclosures attached.

blocked signal. A programmatic signal that is currently in the pending signal mask of a process and, when generated, is not delivered to the process because of the signal mask setting. Some signals cannot be blocked.

block special file. In the Open System Services (OSS) environment, a device that is treated as a file for which all input or output must occur in blocks of data. Traditionally, such files are disk or tape devices. Block special files provide access to a device in a manner that hides the hardware characteristics of the device. Contrast with character special file.

bond. A reliable connection that ensures the required electrical conductivity between conductive parts that must be electrically connected.

bonded. The mechanical interconnection of conductive parts to maintain a common electrical potential.

bonding. The permanent joining of conductive parts to form a path that ensures electrical continuity and the capacity to safely conduct any current likely to be imposed.

bonding jumper. See main bonding jumper.

boot. A synonym for load. Load is the preferred term used in this and other HP NonStop™ S-series system publications.

BOOTP. A protocol for providing initialization information to diskless nodes in an open network.

BOOTPC. See BOOTP client (BOOTPC).

BOOTP client (BOOTPC). A client provided as a Portable Silicon Operating System (pSOS) system product task in the essential firmware on each communications line interface processor (CLIP) in the ServerNet wide area network (SWAN) concentrator. BOOTPC tasks are also provided on the host system as the WANBoot process in the WAN subsystem.

BOOTPD. See BOOTP daemon (BOOTPD).

BOOTP daemon (BOOTPD). The BOOTP server. One BOOTPD runs as a Portable Silicon Operating System (pSOS) system product task in the essential firmware on each communications line interface processor (CLIP) in the ServerNet wide area network (SWAN) concentrator. BOOTPD tasks are also provided on the host system as the WANBoot process in the WAN subsystem.

branch circuit. The circuit conductors located between the equipment receptacles and the final overcurrent device in a power distribution panel (PDP) that protect the circuits.

branded product. A software product that is licensed by X/Open to carry the X/Open or UNIX trademark.

branding process. The activities that lead to the acceptance of a product by X/Open in accordance with its Trade Mark Licence Agreement.

break condition. An event indicator or sequence of data from a terminal or terminal emulator that requests interruption of an application program.

bridge rectifier. A full-wave rectifier with four elements, as in a bridge circuit. Alternating voltage is applied to one pair of opposite junctions, and direct voltage is obtained from the other pair of junctions.

BSD. Berkeley Software Distribution.

built configuration. A configuration revision for which a system image and activation package have been created.

built-in command. In the Open System Services (OSS) environment, a command that is implemented within the `/bin/sh` file. Some built-in commands are also available as separately executable files.

bypass mechanism. Equipment that permits switching from one power source to another. For example, a bypass mechanism on an uninterruptible power supply (UPS) would switch to an alternative power source (such as a standby power generator or commercial utility source) when maintenance must be performed on the UPS.

byte-synchronous. A type of Open Systems Interconnection (OSI) Layer-2 protocol that uses synchronous transmission techniques and requires a character code to define terminal and line control sequences. Data is always transmitted in a block.

cabinet. One or more modules of a system, housed together.

cable channel. A cable management conduit that protects the cables that run between two system enclosures in a double-high stack. Each system enclosure has two cable channels running vertically on its service side: one on the left-hand side of the enclosure, and one on the right-hand side of the enclosure.

cable guidepost. A cable management rod that routes cables exiting the upper enclosure in a double-high stack to prevent the cables from hanging down in front of the customer-replaceable units (CRUs) in the base enclosure. A cable guidepost extends from the base of each cable channel.

cable support. A piece of cable management hardware that secures system cables. The cable support attaches to the service side of a system enclosure near the bottom of the enclosure. Cable ties for securing system cables are threaded through the cable support. The cable support also contains the group and module ID labels and the rear group service light-emitting diode (LED).

cache (cache memory). A small, fast memory holding recently accessed data designed to speed up subsequent access to the same data. Cache memory is built from faster memory chips than main memory, and it is most often used with process or main memory but also used in network data transfer to maintain a local copy of data.

cached bindings. A copy in virtual memory of the data pages containing symbolic references that were rebound when a loadfile was loaded. The cached bindings are associated with a library import characterization that characterizes the set of loadfiles to which the symbols were bound. If the same file is subsequently loaded in an equivalent environment in the same processor, the cached bindings can be reused. See fastLoad.

CAE. See common applications environment (CAE).

canonical input mode. For an Open System Services (OSS) process, a terminal input mode in which data is not made available to the process until an entire logical line (delimited by a newline, EOF, or EOL character) is entered. This mode is sometimes called line mode or nontransparent mode. Contrast with noncanonical input mode.

CAP. See cartridge access port (CAP).

Carbon Copy. A remote operations software application that enables a workstation in one location to access, through a modem, a workstation in another location. Carbon Copy is included with all system consoles, and service providers use it to dial in to system consoles at customer sites. See also remote access.

card cage. A structure made up of slots that hold components such as customer-replaceable units (CRUs) and ServerNet adapters.

carrier. (1) A sheet-metal structure that allows a single-high ServerNet adapter to be installed in a ServerNet adapter slot designed for a double-high ServerNet adapter. (2) An electrical signal that carries data.

cartridge access port (CAP). The component on the tape libraries supported on HP NonStop™ S-series systems where you insert cartridges into and remove cartridges from the library.

caught signal. A programmatic signal that is delivered to a process that has a signal-handling function for it. When the signal is caught, the process is interrupted, and the signal-handling function executes.

CBB. See common base board (CBB).

CCITT. International Telegraph and Telephone Consultative Committee.

CCSA. See Common Communication ServerNet adapter (CCSA).

CE. Customer engineer. See service provider.

cell. See storage pool.

central processing unit (CPU). Historically, the main data processing unit of a computer. The HP NonStop™ servers have multiple cooperating processors rather than a single CPU. See also processor.

Challenge Handshake Authentication Protocol (CHAP). An Internet-standard protocol for verifying encrypted passwords. CHAP is a security protocol that is implemented using Point-to-Point Protocol (PPP). The OSM and TSM Notification Director use CHAP to maintain security during dial-outs.

channel. An information route for data transmission. See also ServerNet link.

CHAP. See Challenge Handshake Authentication Protocol (CHAP).

character. A sequence of one or more bytes representing a single character; used for the organization, representation, or control of data. A single-byte character consists of eight bits that represent a character. A multibyte character uses one or more bytes to represent a character. A wide character is a fixed-width character wide enough to hold any coded character supported by an implementation.

The ISO C standard defines the term multibyte character; a single-byte character is a special case of multibyte character.

character set. A finite set of characters (letters, digits, symbols, ideographs, or control functions) used for the organization, representation, or control of data. See also code set.

character special file. In the Open System Services (OSS) environment, a device that is treated as a file for which all input or output must occur in character bytes. Traditionally, such files are interactive terminals, and the ISO/IEC IS 9945-1:1990 standard defines only the access to such terminal files. See also terminal. Contrast with block special file.

chassis. A single sheet-metal structure that houses one set of system components. In an HP NonStop S-series server, a chassis is part of a system enclosure but can also be mounted in any standard 19-inch rack.

checksum. A generic term, meaning to *add* together (although the definition of *add* need not be a normal arithmetic add) all of the data to produce a check *word*. See also cyclic redundancy check (CRC).

child process. A process created by another process. The creating process becomes the parent process of the new process. See also parent process.

CIIN. A command file in the SYS*nn* subvolume that is read and executed by the startup HP Tandem Advanced Command Language (TACL) process after system load if the CIIN file is specified in the CONFTEXT file and enabled in the OSM or TSM Low-Level Link.

circuit breaker. A device designed to open and close a circuit by nonautomatic means and to open the circuit automatically on a predetermined overcurrent without damage to itself.

CISC. See complex instruction-set computing (CISC).

CISC processor. An instruction processing unit (IPU) that is based on complex instruction-set computing (CISC) architecture.

class. A group of object-oriented data entities and the methods associated with that group.

Class-1 CRU. A customer-replaceable unit (CRU) that probably will not cause a partial or total system outage if the documented replacement procedure is not followed correctly. Customers replacing Class-1 CRUs are not required to have previous experience replacing HP NonStop™ S-series CRUs. However, for some CRUs, customers must be able to use the tools needed for the replacement procedure (which are common tools) and must protect components from electrostatic discharge (ESD).

Class-2 CRU. A customer-replaceable unit (CRU) that might cause a partial or total system outage if the documented replacement procedure is not followed correctly. Customers replacing Class-2 CRUs should have either three or more months of experience replacing HP NonStop™ S-series CRUs or equivalent training. Customers must be able to use the tools needed for the replacement procedure and must protect components from electrostatic discharge (ESD).

Class-3 CRU. A customer-replaceable unit (CRU) that probably will cause a partial or total system outage if the documented replacement procedure is not followed correctly. Customers replacing Class-3 CRUs should have either six or more months of experience replacing HP NonStop™ S-series CRUs or equivalent training. Customers must be able to use the tools needed for the replacement procedure, must protect components from electrostatic discharge (ESD), and must understand the dependencies involved in NonStop S-series CRU-replacement procedures, such as disk-path switching. Replacement by a service provider trained by HP is recommended.

client. A software process, hardware device, or combination of the two that requests services from a server. Often, the client is a process residing on a programmable workstation and is the part of an application that provides the user interface. The workstation client might also perform other portions of the application logic.

client application. An application that requests a service from a server application. Execution of remote procedure calls is an example of a client application.

client (of a loadable library). A loadfile that uses functions or data from a library.

CLIP. See communications line interface processor (CLIP).

cluster. (1) A collection of servers, or nodes, that can function either independently or collectively as a processing unit. See also ServerNet cluster. (2) A term used to describe a system in a Fiber Optic Extension (FOX) ring. More specifically, a FOX cluster is a collection of processors and I/O devices functioning as a logical group. In FOX nomenclature, the term is synonymous with system or node.

cluster number. A number that uniquely identifies a node in a Fiber Optic Extension (FOX) ring. This number is in the range 1 through 14. See also node number.

cluster switch. See HP NonStop™ Cluster Switch (model 6770) and HP NonStop™ ServerNet Switch (model 6780).

cluster switch enclosure. An enclosure provided by HP for housing the subcomponents of an HP NonStop™ Cluster Switch. The subcomponents include the ServerNet II Switch, the AC transfer switch, and the uninterruptible power supply (UPS). A cluster switch enclosure resembles, but is half the height of, a standard HP NonStop S-series system enclosure.

cluster switch group. Within an external ServerNet fabric, all the cluster switches that belong to the same cluster switch zone. A cluster switch group can consist of up to four 6780 switches, each representing one cluster switch layer. All of the cluster switches that form a cluster switch group typically are installed in the same cluster switch rack.

cluster switch layer. The topological cluster switch position within a cluster switch group. Each cluster switch group can contain up to four layers, numbered 1 to 4 from bottom to top. A cluster switch layer consists of a pair of cluster switches (X and Y) and provides connections for up to eight ServerNet nodes. Layers within a group are interconnected by intragroup cables. When all four layers are present, the intragroup cables are configured as a vertical tetrahedron. See also cluster switch layer number.

cluster switch layer number. A number in the range 1 through 4 that identifies the position of a cluster switch within a cluster switch group. See also cluster switch group.

cluster switch logic board. A circuit board that provides switching logic for the HP NonStop™ ServerNet Switch (model 6780). The logic board (LB) has a front panel for operator and maintenance functions and is a Class-3 CRU.

cluster switch rack. A mechanical frame consisting of or based on a 19-inch rack that supports the hardware necessary for a cluster switch group.

cluster switch zone. A pair of X-fabric and Y-fabric cluster switch groups and the ServerNet nodes connected to them. Up to three zones are possible. The zones, if more than one, are interconnected by interzone cables, with each cluster switch layer cabled separately from the other layers.

CME. See correctable memory error (CME).

CMI. See Communications Management Interface (CMI).

code segment. A segment that contains executable instructions of a program or library to be executed plus related information. Code segments can be executed and also accessed as read-only data but not written to by an application program. These read-only and execute-only segments are efficiently shared among simultaneous executions of that program or library. Therefore, they are read from disk but are never written back to disk. See also TNS code space.

code set. Codes that map a unique numeric value to each character in a character set, using a designated number of bits to represent each character. Single-byte code sets use 7 or 8 bits to represent each character. The ASCII and ISO 646 code sets use 7 bits to represent each character in Roman-based alphabets; these code sets are very limited and are not appropriate for international use. The single-byte ISO 8859 code sets use 8 bits to represent each character and can therefore support Roman-based alphabets and many others including Greek, Arabic, Hebrew, and Turkish. Multibyte code sets represent characters that require more than one byte, such as East Asian ideographic characters.

code space. See TNS code space.

cold load. A synonym for system load or load (in the case of single processor load). System load or load is the preferred term in HP NonStop™ S-series system publications.

command. A demand for action by or information from a subsystem or the operation demanded by an operator or application. A command is typically conveyed as an interprocess message from an application to a subsystem.

command file. An EDIT file that contains a series of commands and serves as a source of command input.

common applications environment (CAE). A computer environment in which applications can be ported across all X/Open branded products because of the use of international and industry standards. A CAE is an open system application development environment, an open system execution environment, or a combination of the two.

common base board (CBB). In modular customer-replaceable units (CRUs), the printed wiring assembly (PWA) that plug-in cards (PICs) are installed in.

Common Communication ServerNet adapter (CCSA). A ServerNet adapter that provides an HP NonStop™ S-series integration platform for Signaling System Number 7 (SS7) protocol communications.

common mode. Electrical interference that can be measured as a ground-referenced signal. In true common mode, a signal is common to all of the current-carrying conductors.

common-mode transients. Transients that appear between both inputs of a circuit and a common reference (such as ground).

communications line. A two-way link consisting of processing equipment, I/O devices, protocol conventions, and cables that connect a computer to other computers.

communications line interface processor (CLIP). The major programmable device within the ServerNet wide area network (SWAN) concentrator, providing link-level protocol and a software interface to the host. The CLIP stores and implements specific communications protocols.

Communications Management Interface (CMI). A utility used in D-series and earlier release version updates (RVUs) to make online changes to the configuration of ATP6100, CP6100, and EnvoyACP/XF communications subdevices. In G-series RVUs, CMI functions are performed by the Subsystem Control Facility (SCF).

communications subsystem. The combination of data communications hardware and software processes that function together as an integrated unit to provide services and access to wide and local area networks.

Compaq TSM. Identifies a client or server software component used to manage or service HP NonStop™ S-series servers. See also Compaq TSM client software and Compaq TSM server software.

Compaq TSM client software. The component of the Compaq TSM package that runs on a system console. The TSM client software consists of the TSM Low-Level Link, the TSM Service Application, the TSM Notification Director, and the TSM Event Viewer. See also Compaq TSM server software.

Compaq TSM Event Viewer. A component of the Compaq TSM client software. The TSM Event Viewer lets you set up criteria to view Event Management Service (EMS) log files in several ways, enabling you to rapidly assess service problems.

Compaq TSM Low-Level Link. A component of the Compaq TSM client software. The TSM Low-Level Link enables you to communicate with an HP NonStop™ S-series server even when the HP NonStop Kernel operating system is not running. When the operating system is running, you usually communicate with the server using the TSM Service Application. See also Compaq TSM Service Application.

Compaq TSM package. A software product for HP NonStop™ S-series servers that provides the information needed to perform functions such as querying resources and testing, provides notification of problems on the system, and allows local or remote access to the system for service and maintenance. The TSM package performs the same role as that of HP Tandem Maintenance and Diagnostic System (TMDS), Syshealth, and Remote Maintenance Interface (RMI) on earlier systems.

Compaq TSM server software. The component of the Compaq TSM package that runs on an HP NonStop™ S-series server. When the HP NonStop Kernel operating system is running, the TSM client software on a system console communicates with a server through the TSM server software. See also Compaq TSM client software.

Compaq TSM Service Application. A component of the Compaq TSM client software. The TSM Service Application enables you to communicate with an HP NonStop™ S-series server when the HP NonStop Kernel operating system is running. When the operating system is not running, communication must take place using the TSM Low-Level Link. See also Compaq TSM Low-Level Link.

compiler extended-data segment. A selectable segment, with ID 1024, created and selected automatically in many (but not all) TNS processes. Within this segment, the compiler automatically allocates global and local variables and heaps that would not fit in the TNS user data segment. A programmer must keep this segment selected whenever those items might be referenced. Any alternative selections of segments must be temporary and undone before returning.

complex instruction-set computing (CISC). A processor architecture based on a large instruction set, characterized by numerous addressing modes, multicycle machine instructions, and many special-purpose instructions. Contrast with reduced instruction-set computing (RISC).

compliance. The testing and verification process that precedes X/Open licensing.

computer-room power center (CRPC). The equipment that conditions and distributes facility power to computer-room equipment. The CRPC typically houses an electrostatically shielded isolation transformer, power distribution panels (PDPs), a main shunt-trip circuit breaker, and voltage indicators. Also referred to as a power distribution unit (PDU) or power distribution center.

concentrator manager process (ConMgr). A process provided as part of the wide area network (WAN) subsystem. The ConMgr process runs in each processor that supports WAN products and provides management functions to the WAN subsystem and WAN products, such as downloading data link control (DLC) tasks to the communications line interface processors (CLIPs) on the ServerNet wide area network (SWAN) concentrator and selecting the preferred path for the DLC tasks.

conduit. A tubular raceway, usually constructed of rigid or flexible metal, through which insulated power and ground conductors or data cables are run. Nonmetallic conduits, although available, are not recommended.

CONFAUX file. The auxiliary configuration file created by the Distributed Systems Management/Software Configuration Manager (DSM/SCM) tools. The CONFAUX file contains a list of the code files and system files that are needed to build the new HP NonStop™ Kernel operating system. HP recommends that you avoid making any changes to your CONFAUX file.

CONFBASE file. In G-series release version updates (RVUs), the basic system configuration database file, which is stored on the $SYSTEM.SYS*nn* subvolume. See also configuration file.

CONFIG file. In G-series release version updates (RVUs), the current system configuration database file, which is stored on the $SYSTEM. ZSYSCONF subvolume. See also configuration file.

configuration. (1) The arrangement of enclosures, system components, and peripheral devices into a working unit. (2) The definition or alteration of characteristics of an object.

configuration file. In G-series release version updates (RVUs), one of the following files: CONFBASE, CONFIG, one or more saved configuration files named CONF*xxyy*, and CONFSAVE. See also system configuration database. In RVUs prior to the G-series, the configuration file is either the OSCONFIG file used by the Configuration Utility Program (COUP) or the CONFTEXT file used during system generation.

configuration planner. The person who manages system configuration changes and software configuration changes. This person modifies the system configuration database for system configuration changes and creates a new operating system image for software configuration changes. See also planner.

configuration revision. A planner-defined set of software products and related configuration information that the Distributed Systems Management/Software Configuration Manager (DSM/SCM) can activate on a target system. Multiple configuration revisions might exist on a target system. A configuration revision is made up of the product versions named in its software revision list, its HP NonStop™ Kernel operating system image, and the relevant profile items, such as the location of the target subvolumes on the target system. It is created by a Build request and is included in the activation package sent to the target system.

configuration tag. Each configuration tag identifies and configures the topology of a cluster switch and its unique position within the topology. The configuration tag defines which ServerNet node numbers the cluster switch supports.

configuration utility process. The $ZCNF process that is the access process for the CONFIG file and starts and maintains the $ZPM persistence manager process.

Configuration Utility Program (COUP). A utility used in D-series and earlier release version updates (RVUs) to make online changes to the configuration of devices and controllers. COUP is part of the Dynamic System Configuration (DSC) facility. In G-series RVUs, similar functions are performed by the Subsystem Control Facility (SCF).

configured object. A Subsystem Control Facility (SCF) object that exists at the time a subsystem completes its initialization process, or an SCF object that is brought into existence by a command issued through a subsystem management interface.

CONFLIST file. The output file produced during system generation, including error and warning messages.

conformance. Meeting the requirements of a specific standard.

conformance document. An implementor’s document that must accompany software claiming conformance with a POSIX standard. The document specifies the behavior or other aspect of the software when the standard describes a behavior or aspect as implementation-defined.

conformance statement questionnaire (CSQ). A document that identifies how a product implements X/Open Specifications as defined in XPG Component/Profile Definitions. A CSQ exists for each branded product.

conforming POSIX.1 application. An application that is either an ISO/IEC-conforming POSIX.1 application or a national-standards-body conforming POSIX.1 application.

conforming POSIX.1 application using extensions. An application that:

- Is a conforming POSIX.1 application
- Also uses features or facilities that are not described in ISO/IEC IS 9945-1:1990 (POSIX.1) but are consistent with the standard
- Meets the documentation requirements of a conforming POSIX.1 application and documents its use of nonstandard features or facilities

For example, an application using the `tdm_fork()` function could be a conforming POSIX.1 application using extensions.

CONFSAVE file. In G-series release version updates (RVUs), the automatically saved configuration database file, which is stored on the $SYSTEM.ZSYSCONF subvolume. See also configuration file.

CONFTEXT file. The configuration file used as input during system generation that contains a series of entries defining your HP NonStop™ Kernel operating system attributes. A G-series CONFTEXT file consists of one or two paragraphs: DEFINES (optional) and ALLPROCESSORS.

CONF*xxyy* file. In G-series release version updates (RVUs), a saved configuration database file created by the Subsystem Control Facility (SCF) and stored in the $SYSTEM.ZSYSCONF subvolume. *xxyy* is the number you entered as *xx.yy* in the SCF SAVE CONFIGURATION command. (*xx* indicates the base version, and *yy* indicates the subversion.) See also configuration file.

ConMgr. See concentrator manager process (ConMgr).

connection. (1) The path between two protocol modules that provides reliable stream delivery service. (2) For OSM and TSM software, the logical link established between the client software on a workstation and the server software on an HP NonStop™ S-series system after a logon sequence has been performed. The two types of logical connections are service connections and low-level links.

Connection view. One of several views of a server available in the view pane of the Management window of the Compaq TSM package. The Connection view is a visual representation of the connectivity among components within an enclosure. See also Physical view.

connectivity. The ability of a system to transfer information between itself and a system from another vendor. Other vendors use the term *connectivity* to mean hardware compatibility. See also interoperability.

connector. See port.

console message. See operator message.

contiguous ground. An insulated grounding conductor that extends from an equipment enclosure power receptacle to the final point of electrical service for the computer-room equipment, whether that final point is the main service entrance or the separately derived power source. In most instances, the final point of electrical service is an isolating transformer installed in the computer room.

control and inquiry. The aspects of Subsystem Control Facility (SCF) object management related to the state or configuration of an object. Such aspects include actions that affect the state or configuration of an object, inquiries about the object, and commands pertaining to the session environment (for example, commands that set default values for the session).

controller. See I/O controller or ServerNet addressable controller (SAC).

controlling process. In the Open System Services (OSS) environment, the session leader that established the connection to the controlling terminal. The session leader stops being the controlling process when the corresponding terminal stops being the controlling terminal.

controlling terminal. In the Open System Services (OSS) environment, a terminal that might be associated with a session. A session can have only one controlling terminal, and a controlling terminal can control only one session at a time. When a session has a controlling terminal, all the following are true:

- Certain character sequences entered from that terminal cause signals to be sent to all processes in the process groups of that session.
- Certain characters entered from that terminal might receive special treatment.
- Members of background process groups of the session are restricted from certain kinds of access to the controlling terminal.

Coordinated Universal Time (UTC). The standard measure of time from the beginning of the current Epoch. UTC is sometimes called Universal Coordinated Time, CUT, or UCT; the standard appellation is abbreviated as UTC, an arbitrary ordering of the letters. UTC was formerly called Greenwich mean time (GMT).

core dump file. See saveabend file.

core file. See saveabend file.

correctable memory error (CME). An error caused by incorrect data at a particular memory location. The cause of the error is such that the error is automatically corrected by the system. Contrast with uncorrectable memory error (UCME).

COUP. See Configuration Utility Program (COUP).

CPU. See central processing unit (CPU).

cpu, pin. In the Guardian environment, a number pair that uniquely identifies a process during the lifetime of the process, consisting of the processor (CPU) number and the process identification number (PIN). See also PID.

CRC. See cyclic redundancy check (CRC).

creation version serial number (CRVSN). In the Open System Services (OSS) environment, a number assigned by a disk process when a file is created. The CRVSN is used by the disk process and the OSS name server process to verify that the correct file is accessed. The CRVSN is stored in the catalog entry for an OSS regular file and is passed to the disk process when a Data Definition Language (DDL) request that involves the file is made.

critical load. Equipment that must have an uninterruptible power input to prevent damage to the equipment or the facility or prevent injury to personnel.

CRPC. See computer-room power center (CRPC).

CRU. See customer-replaceable unit (CRU).

CRVSN. See creation version serial number (CRVSN).

CSQ. See conformance statement questionnaire (CSQ).

current. The movement of electrons caused by potential difference between two electromotive charge forces.

current configuration file. See configuration file.

current working directory. In the Open System Services (OSS) environment, the directory used in pathname resolution of relative pathnames. A process always has a current working directory. See also working directory.

CUSTFILE. An EDIT file included on every site update tape (SUT) as $SYSTEM.A*nnnnnn*.CUSTFILE, where *nnnnnn* is the system serial number of the target system. The CUSTFILE contains information on the software products on the SUT, their related files, and the destination and use of each file. HP customizes information in the CUSTFILE for each customer’s system.

customer engineer (CE). See service provider.

customer-installable system. A system that does not require specially trained service providers to install.

customer-replaceable unit (CRU). A unit that can be replaced in the field either by customers or by qualified personnel trained by HP. CRUs are divided into the categories of Class 1, Class 2, and Class 3 based on the risk of causing a system outage if the documented replacement procedure is not followed correctly and how much CRU-replacement training or experience is advisable. See also Class-1 CRU, Class-2 CRU, Class-3 CRU, and field-replaceable unit (FRU).

cyclic redundancy check (CRC). The most widely used error detection code for ensuring the integrity of transmitted data. The digits of the CRC are calculated by the sender for each block of data sent and recalculated by the receiver. (It is a family of mathematical functions involving computing the quotient and remainder of a polynomial division.) A CRC is a form of checksum.

daemon. See demon.

dark site. See unattended site.

data communications equipment (DCE). Equipment that provides all the functions required to establish, maintain, and terminate a connection and provides the signal conversion and coding between the data terminal equipment (DTE) and telephone company lines or data circuits. A DCE is usually a modem.

data link control (DLC). A set of functions associated with Layer 2 of the Open Systems Interconnection (OSI) reference model. These functions are responsible for reliable communication between two physically connected nodes.

data link control (DLC) task. Tasks that support the equivalent to Layer 2 of the Open Systems Interconnection (OSI) reference model. Wide area network (WAN) DLC tasks execute in the ServerNet wide area network (SWAN) concentrator communications line interface processor (CLIP), and each WAN DLC task controls one line interface.

data segment. A virtual memory segment holding data. Every process begins with its own data segments for program global variables and runtime stacks (and for some libraries, instance data). Additional data segments can be dynamically created. See also flat segment and selectable segment.

data terminal equipment (DTE). Equipment that constitutes the data source or data sink and provides for the communication control function protocol. It includes any piece of equipment at which a communication path begins or ends.

data transparent. Describes software that examines all eight bits of every data byte and that uses no bit in a data byte for its own purposes. Internationalized applications must be data transparent.

dB. See decibel (dB).

dBm. Decibels as referenced to a milliwatt. A unit of measure that establishes 0 dBm equal to 1 milliwatt. A negative value represents a decrease in power, and a positive value represents an increase in power. See also decibel (dB).

DC. See direct current (DC).

DCE. See data communications equipment (DCE).

DCF. See dynamic configuration file (DCF).

DCT. See destination control table (DCT).

DC power cable. In system enclosures with power shelves, a cable that delivers DC power from the power shelf to a processor multifunction (PMF) customer-replaceable unit (CRU) or I/O multifunction (IOMF) CRU in that enclosure.

decibel (dB). A unit of measure used to express a relative difference in power. A negative value represents a decrease in power, and a positive value represents an increase in power.

dedicated service LAN. An Ethernet local area network (LAN) for use by only OSM and TSM applications. This LAN connects system consoles with the Ethernet ports on the processor multifunction (PMF) customer-replaceable units (CRUs) in group 01 of an HP NonStop™ S-series server. A dedicated LAN supports NonStop S-series servers and system consoles but does not support any other types of servers or workstations. See also public LAN.

DEFINE. An HP Tandem Advanced Command Language (TACL) command you can use to specify a named set of attributes and values to pass to a process.

DEFINES paragraph. An optional paragraph in the CONFTEXT configuration file that contains one or more identifiers, each with its associated text string. The DEFINES paragraph, if used, precedes the ALLPROCESSORS paragraph.

delta. A method for connecting a 3-phase power source (or load) in a closed series loop with input (or output) connections made to each of the three junctions. The delta's physical arrangement resembles the delta character from the Greek alphabet.

demon. On a UNIX system, a process that runs continuously to provide a specific service for other processes. A demon does not have a controlling terminal and is not explicitly invoked. On an HP NonStop™ system, a demon runs in the Open System Services (OSS) environment and has an OSS process ID. See also static server.

destination control table (DCT). The HP NonStop™ Kernel data structure that holds information about every device and named process in the system. The DCT consists of the named resource list (NRL) and the process-pair list (PPL). The DCT is replicated in each CPU.

destination ServerNet ID (DID). A field in the ServerNet packet header indicating the intended destination for the packet.

detailed report. A complete listing of status or configuration information provided by the Subsystem Control Facility (SCF) STATUS or INFO command when you use the DETAIL option. Contrast with summary report.

device. A computer peripheral or an object that appears to an application as such. See also terminal.

DHCP. See Dynamic Host Configuration Protocol (DHCP).

dial-out point. A system console from which incident reports are sent to a service provider. Incident reports are sent only from system consoles defined as the primary and backup dial-out points (the primary and backup system consoles).

DID. See destination ServerNet ID (DID).

DIMM. See dual inline memory module (DIMM).

direct connect. Connection from the Fibre Channel ServerNet adapter (FCSA) to the Enterprise Storage System (ESS) without going through an FC switch.

direct current (DC). Electric current that flows in only one direction. Contrast with alternating current (AC).

direct jump area. One of sixteen 256-megabyte portions of the 4-gigabyte virtual address space. A RISC jump instruction has the ability to jump directly to any location within its own direct jump area without having to use a far jump table.

directory. A type of Open System Services (OSS) special file that contains directory entries, which associate names with files. No two directory entries in the same directory have the same name.

directory entry. In the Open System Services (OSS) file system, an object that associates a filename with a file. Several directory entries can associate names with the same file. See also link.

directory loop. In the Open System Services (OSS) file system, an error condition in which a directory is identified as its own parent directory.

directory special file. See directory.

directory stream. In the Open System Services (OSS) file system, an object with an opaque data type. A process can sequentially read directory entries from a directory stream.

directory tree. A hierarchy of directories. In the Open System Services (OSS) environment, directories are connected to each other in a branching hierarchical fashion such that only one path exists between any two directories (if no backtracking occurs).

disconnecting means. A device, group of devices, or other means by which the conductors of a circuit can be disconnected from their source of supply.

discovery. For the OSM and TSM client software, the process of identifying the resources that exist on an HP NonStop™ S-series server. See also incremental discovery and initial discovery.

disk bootstrap. A software entity residing on disk that is used to load the HP NonStop™ Kernel operating system image (OSIMAGE) into memory during a system load. A disk that contains the disk bootstrap is referred to as a bootable disk. The disk bootstrap is placed on the disk either as part of a tape load or as a result of the SCF CONTROL DISK, REPLACEBOOT command.

disk cache. A temporary storage buffer into which data is read, retained, and perhaps updated before being written to disk, for more efficient processing.

disk drive. A device that stores and accesses data on a magnetic disk. Random access to addressable locations on a magnetic disk is provided by magnetic read/write heads. See also volume.

DISKGEN. A system generation option that invokes the DISKGEN program to copy directly to disk the files necessary to generate an HP NonStop™ Kernel operating system. DISKGEN can be used instead of a system image tape (SIT).

DISK object type. The Subsystem Control Facility (SCF) object type for all disk devices attached to your system.

disk volume. See volume.

distributed system. A system that consists of a group of connected, cooperating computers.

Distributed Systems Management (DSM). A set of tools used to manage HP NonStop™ S-series systems and Expand networks.

Distributed Systems Management/Software Configuration Manager (DSM/SCM). A graphical user interface (GUI)-based program that installs new software and creates a new HP NonStop™ Kernel operating system. DSM/SCM creates a new software revision and activates the new software on the target system.

distribution subvolume (DSV). A subvolume containing program files for a particular software product along with the release version update (RVU) document (softdoc) file for that product. The format for a DSV name is Y*nnnnrrr* or R*nnnnrrr*, where *nnnn* is the software product number and *rrr* is the base version identifier (such as D20) or software product revision (SPR) identifier (such as AAB).

DLC. See data link control (DLC).

DNS. See Domain Name System (DNS).

DNS server. A server that resolves hostnames to Internet protocol (IP) address mapping queries. These queries originate from either client computers, which are known as resolvers, or other Domain Name System (DNS) servers, which accounts for the distributed nature of DNS. See also Network Information Service (NIS).

domain. (1) In the Internet, a part of the naming hierarchy. Syntactically, a domain name consists of a sequence of names (labels) separated by periods (dots).

(2) In a NonStop S-series server, a pair of service processors, the associated router clouds, and the attached replaceable units.

(3) A set of objects over which control or ownership is maintained. Types of domains include power domains and service processor (SP) domains.

Domain Name System (DNS). A system that defines a hierarchical, yet distributed, database of information about hosts on a network. The network administrator configures the DNS with a list of hostnames and Internet protocol (IP) addresses, allowing users of workstations that are configured to query the DNS to specify remote systems by hostnames rather than by IP addresses. DNS domains should not be confused with Windows NT networking domains. See also DNS server, Network Information Service (NIS), and ping.

donor system. The computer system you make smaller by removing enclosures, either to reduce the system or to add the removed enclosures to another target system, using a process known as system reduction.

double-high ServerNet adapter. A ServerNet adapter that occupies an entire ServerNet adapter slot in an HP NonStop™ S-series server. Contrast with single-high ServerNet adapter.

double-high stack. A stack that includes a base, a frame, and two system enclosures. Contrast with single-high stack.

double-wide plug-in card (PIC). A large-form-factor plug-in card (PIC) that occupies two adjacent PIC slots within a customer-replaceable unit (CRU). See also single-wide plug-in card (PIC).

download. The process of transferring software from one location to another, where the transferring entity initiates the transfer.

download line task. Any task running under the Portable Silicon Operating System (pSOS) system product, such as a data protocol.

downtime. Time during which a computer system is not capable of doing useful work because of a planned or unplanned outage. From the end user’s perspective, downtime is any time a needed application is not available.

downward compatibility. The ability of a requester to operate with a server of an earlier revision level. In this case, the requester is downward-compatible with the server, and the server is upward-compatible with the requester. Contrast with upward compatibility.

drive. See disk drive or tape drive.

dropout. A voltage loss of very short duration (that is, milliseconds).

DSC. See Dynamic System Configuration (DSC).

DSM. See Distributed Systems Management (DSM).

DSM/SCM. See Distributed Systems Management/Software Configuration Manager (DSM/SCM).

DSV. See distribution subvolume (DSV).

DTE. See data terminal equipment (DTE).

dual inline memory module (DIMM). Small circuit boards carrying memory integrated circuits, with signal and power pins on both sides of the board. A DIMM is different from a single inline memory module (SIMM) in that the connections on each side of the module connect to different chips, whereas the connections on both sides of a SIMM connect to the same memory chip. This difference gives the DIMM a wider data path, as more modules can be accessed at once.

dual-ported. The capability of a ServerNet adapter or peripheral device to receive data and commands from two sources although only one source might have access at any particular moment.

duplicate file descriptor. In the Open System Services (OSS) file system, a file descriptor that refers to the same open file description as another file descriptor.

dynamic configuration file (DCF). An attachment file that is produced by the OSM and TSM client software and accompanied by an incident report. The DCF contains a snapshot of the system configuration, the state of the HP NonStop™ S-series server, and outstanding alarms at the time that the incident report was issued. The DCF is used by the service provider to avoid having to perform online discovery of the server over dial-up telephone lines.

Dynamic Host Configuration Protocol (DHCP). An Internet protocol for automating the configuration of computers that use TCP/IP. DHCP can automatically assign IP addresses, deliver TCP/IP stack configuration parameters such as the subnet mask and default router, and provide other configuration information such as the addresses for printer, time, and news servers.

dynamic information. Information that represents the set of resources that actually exist in the current configuration of an HP NonStop™ S-series server. Dynamic information is gathered from a server through the process of discovery. Contrast with static information.

dynamic-link library (DLL). A collection of procedures whose code and data can be loaded and executed at any virtual memory address, with run-time resolution of links to and from the main program and other independent libraries. The same DLL can be used by more than one process. Each process gets its own copy of DLL static data. Contrast with shared run-time library (SRL). See also TNS/R library.

dynamic loading. Loading and opening dynamic-link libraries under programmatic control after the program is loaded and execution has begun.

dynamic process configuration. Using Subsystem Control Facility (SCF) to configure a generic process to always start in a designated primary processor (that is, to be fault tolerant).

dynamic shared object (DSO). See dynamic-link library (DLL).

Dynamic System Configuration (DSC). A utility used in D-series and earlier release version updates (RVUs) to make online changes to the configuration of devices and controllers. Its interactive utility is called the Configuration Utility Program (COUP). In G-series RVUs, similar functions are performed by Subsystem Control Facility (SCF).

E4SA. See Ethernet 4 ServerNet adapter (E4SA).

earth ground. The connection of the electrical-grounding conductors to a dependable, low-resistance contact with the soil.

earth-grounding electrode. An electrically conductive rod that is driven into soil, thus providing an earth-ground connection point for the electrical ground wiring in a building. A vertical steel column of a building, with its base sunk into soil, can also serve as an earth-grounding electrode.

earth-grounding electrode system. A grounding network created by bonding together the grounding means in a building (for example, underground metal water pipes, structural steel, and ground rods into the earth) and bonding them to the switchgear at the facility's main electrical service entrance.

ECL. See emitter-coupled logic (ECL).

ECL plug-in card (PIC). See emitter-coupled logic (ECL) plug-in card (PIC).

ECL ServerNet cable. See emitter-coupled logic (ECL) ServerNet cable.

EDIT file. In the Guardian file system, an unstructured file with file code 101. An EDIT file can be processed by either the EDIT or PS Text Edit (TEDIT) editor. An EDIT file typically contains source program or script code, documentation, or program output. Open System Services (OSS) functions can open an EDIT file only for reading.

effective group ID. An attribute of an Open System Services (OSS) process that is used to determine permissions such as the file access allowed for the process. The effective group ID of a process is a group ID that contributes to the group access privileges of that process. The effective group ID of a process might be used to set the group ID of files created by that process. The effective group ID can be changed while the process runs.

effective user ID. An attribute of an Open System Services (OSS) process that is used to determine such permissions as the file access allowed for the process. The effective user ID of a process is the user ID that determines the owner access privileges of that process. The effective user ID of a process might be used to set the user ID of files created by that process. The effective user ID can be changed while the process runs.

EIA. Electronic Industries Association.

electric utility. The local utility service that, for a fee, supplies alternating-current (AC) power to businesses and residences.

electromagnetic interference (EMI). Forms of conducted or radiated interference that might appear in a facility as either normal or common-mode signals. The frequency of the interference can range from the kilohertz to gigahertz range. However, the most troublesome interference signals are usually found in the kilohertz to low megahertz range. At present, the terms electromagnetic interference and radio frequency interference (RFI) are usually used interchangeably.

electrostatically shielded transformer. A transformer that has a metallic shield placed between the primary and secondary windings. This shield diverts high-frequency signals to ground.

electrostatic discharge (ESD) protection kit. A kit containing an antistatic mat and a wriststrap with a cable and grounding clip. A service provider or customer wears the wriststrap while performing maintenance procedures inside an enclosure. The wriststrap and cable contain grounding wires. When the grounding clip is attached to a metal object, such as the enclosure, the person wearing the wriststrap is grounded. Any static electricity incurred during the procedure is discharged safely to the enclosure instead of to electrical components within the enclosure.

ELF. See extended link format (ELF).

emergency power off (EPO). Describes equipment used to automatically disconnect all electrical power to connected equipment if an emergency occurs. A computer room’s main EPO system shuts off all room equipment (except for lighting and fire-sensor equipment) if a fire occurs. An equipment zone EPO shuts off power to all connected computer equipment if a power anomaly occurs.

emergency power-off (EPO) connector. A two-pin connector on the service side of an enclosure that allows an external signal to disable the batteries in the enclosure during emergency conditions. A cable is attached from the connector to a relay band or push button typically located near the door of a computer room. Pushing the EPO button removes power from all computer equipment in the room. The EPO connectors prevent the batteries from powering the server after power is removed. EPO capabilities are required in the United States when a server is installed in a computer room designed to comply with the special construction and fire protection provisions of the United States' national electrical code (or at other sites as required by local regulations.)

EMI. See electromagnetic interference (EMI).

emitter-coupled logic (ECL). A logic that expresses digital signals in differential negative voltage levels, from -8 volts to -1.8 volts. HP NonStop™ S-series servers containing ServerNet expansion boards (SEBs) use ECL ServerNet cables. An ECL plug-in card (PIC) allows the modular SEB (MSEB) and I/O multifunction (IOMF) 2 customer-replaceable unit (CRU) to use ECL ServerNet cables.

emitter-coupled logic (ECL) plug-in card (PIC). A plug-in card (PIC) for the modular ServerNet expansion board (MSEB) and I/O multifunction (IOMF) 2 customer-replaceable unit (CRU) that supports the emitter-coupled logic (ECL) interface. See also emitter-coupled logic (ECL) and plug-in card (PIC).

emitter-coupled logic (ECL) ServerNet cable. A ServerNet cable that uses emitter-coupled logic (ECL). Before the modular ServerNet expansion board (MSEB) was introduced, ECL was the only ServerNet cable technology used by HP NonStop™ S-series servers. You can connect an ECL ServerNet cable directly to a ServerNet expansion board (SEB) or to an MSEB using an ECL plug-in card (PIC).

empty directory. In the Open System Services (OSS) file system, a directory that contains only an entry for itself and an entry for its parent directory.

empty string. In C and C++ programs, a character string that begins with a null character. This term is synonymous with *null string*.

EMS. See Event Management Service (EMS).

EMS collector. An Event Management Service (EMS) process to which subsystems report events.

enclosure. Similar to a cabinet in HP NonStop™ K-series systems. An enclosure can contain components of a system or a peripheral. Base enclosures are placed on the floor and can have other enclosures stacked on top of them. Stackable enclosures can be placed on top of other enclosures. See also system enclosure and peripheral enclosure.

enclosure interleaving. On HP NonStop™ S-series systems, configuring a mirrored disk volume to use two separate system enclosures. For internal disk drives, the two disk drives of the mirrored volume can be in separate enclosures. For external disk drives, the adapters connected to the two disk drives of the mirrored volume can be in separate enclosures.

endian. Denotes the significance of byte 0 in a multibyte structure such as a word. NonStop servers are big-endian, where the most significant bit is contained in byte 0. Intel systems and HP AlphaServer OpenVMS and HP AlphaServer Tru64 UNIX systems are little-endian, where the least significant bit is contained in byte 0.

environmental parameters. Subsystem Control Facility (SCF) session parameters set by default or by using various SCF commands. The values associated with the environmental parameters can be examined using the ENV command.

environment strings. For an Open System Services (OSS) process, a vector of strings of the form *name* = *value* that contains information about the environment that the process runs in. Environment strings are accessible to the process and are inherited by its child processes.

Enterprise Storage System (ESS). A collection of magnetic disks, their controllers, and the disk cache in a stand-alone cabinet or cabinets. These disks are configured by an attached console and presented to the attached server as logical volumes that can be a fraction of a physical volume or can span volumes.

EPO. See emergency power off (EPO).

Epoch. The period beginning January 1, 1970, at 0 hours, 0 minutes, and 0 seconds Coordinated Universal Time (UTC).

EPO connector. See emergency power-off (EPO) connector.

equipment grounding conductor. The conductor used to connect the non-current-carrying metal parts of equipment, raceways, and other enclosures to the grounding electrode conductor at the facility’s main service entrance or at the source of a separately derived power source.

errno. An external variable that contains the most recent error condition set by a C function.

error number. For the Subsystem Programmatic Interface (SPI), a value that can be assigned to a return token or to the last field of an error token to identify an error that occurred. SPI defines a small set of error numbers, but most error numbers are defined by subsystems.

ESD kit. See electrostatic discharge (ESD) protection kit.

ESP. See expansion service processor (ESP).

ESS. See Enterprise Storage System (ESS).

essential firmware. Code in memory that is necessary for power-up initialization and communication with a host or device. Contrast with nonessential firmware.

Ethernet. A local area network (LAN) that uses the carrier sense multiple access with collision detection (CSMA/CD) access method on a bus topology and is the basis for the IEEE 802.3 standard.

Ethernet 4 ServerNet adapter (E4SA). A ServerNet adapter for Ethernet local area networks (LANs) that contains four Ethernet ports.

Ethernet hub. A multiport repeater typically supporting 10Base-T cabling. Most hubs are connectors for 8 or 12 cables. Also referred to as a concentrator.

Event Management Service (EMS). A Distributed Systems Management (DSM) product that provides event collection, event logging, and event distribution facilities. EMS provides different event descriptions for interactive and programmatic interfaces, lets an operator or an application select specific event-message data, and allows for flexible distribution of event messages within a system or network.

event message. Text intended for a system operator that describes a change in some condition in the system or network, whether minor or serious. The change of condition is called an event. Events can be operational errors, notifications of limits exceeded, requests for actions needed, and so on. See also operator message.

Event Viewer Server Manager. A persistent process that routes messages and data between the OSM or TSM Event Viewer on the system console and event server and summary server processes on the HP NonStop™ S-series server. An event server process retrieves events. A summary server is a persistent process that maintains a summary of all the events in a specified log file.

exception handler. A section of program code to which control is transferred when an exception occurs. The exception handler then determines what action should be taken.

Expand line-handler process. A process pair that handles incoming and outgoing Expand messages and packets. An Expand line-handler process handles direct links and also binds to other processes using the Network Access Method (NAM) interface to support Expand-over-X.25, Expand-over-FOX, Expand-over-ServerNet, Expand-over-TCP/IP, and Expand-over-SNA links. See also Expand-over-ServerNet line-handler process.

Expand network. The HP NonStop™ Kernel operating system network that extends the concept of fault-tolerant operation to networks of geographically distributed HP NonStop S-series systems. If the network is properly designed, communication paths are constantly available even if a single line failure or component failure occurs.

Expand node. A system in an Expand network. See also node.

Expand node number. A number in the range 0 through 254, sometimes referred to as the system number, that identifies a node in an Expand network. Each Expand node number must be unique within the network. See also ServerNet node number.

Expand-over-ServerNet line. The single line associated with an Expand-over-ServerNet line-handler process. An Expand-over-ServerNet line has the same name and logical device number as its Expand-over-ServerNet line-handler process. However, the line does not have the same states as the line-handler process.

Expand-over-ServerNet line-handler process. An Expand line-handler process that uses the NETNAM protocol to access the Network Access Method (NAM) interface provided by the ServerNet cluster monitor process, $ZZSCL. The Expand-over-ServerNet line-handler process handles incoming and outgoing Expand messages. It also handles packets leaving the server and security-related messages going between systems within a ServerNet cluster. Each node in a ServerNet cluster must be configured with an Expand-over-ServerNet line-handler process for every other node in the ServerNet cluster.

expansion service processor (ESP). A service processor (SP) that is not a master service processor (MSP). ESPs occur in pairs in groups 02 through *nn* (not in group 01). See also master service processor (MSP).

explicit DLL. See explicit dynamic-link library (explicit DLL).

explicit dynamic-link library (explicit DLL). A dynamic-link library (DLL) that is named in the libList of a client or is a native-compiled loadfile associated with a client.

export. To offer a symbol definition for use by other loadfiles. A loadfile exports a symbol definition for use by other loadfiles that need a data item or function having that symbolic name.

export digest. A mathematical hash of the exported symbol names and locations in a library. Two libraries with the same export digest are interchangeable in that they both export the same symbols at the same locations. They are not necessarily semantically equivalent.

extended data segment. See selectable segment.

extended link format (ELF). A standard binary file format common on UNIX systems. The ELF format is used for position-independent code (PIC) files.

extensible input/output (XIO). A redesign of the HP NonStop™ Kernel operating system's I/O subsystem to enable it to extend itself in general ways to meet future requirements.

extent. A contiguous area on disk for allocating one file.

extent size. The size in bytes of a contiguous area on disk for allocating one file.

external entry point (XEP) table. A table located in the last page of each TNS code segment that contains links for calls (unresolved external references) out of that segment.

external fabric connection. The low-level ServerNet connection between a node and one of the external ServerNet fabrics. Each node has an X and a Y connection to the external fabrics.

external routing. The routing of packets over the external ServerNet fabrics; that is, between systems (or nodes) in a ServerNet cluster. See also internal routing.

external ServerNet fabrics. The fabrics that link systems in a ServerNet cluster. See also internal ServerNet fabrics.

external ServerNet X or Y fabric. The X or Y fabric that links systems in a ServerNet cluster. See also internal ServerNet X or Y fabric.

external system area network manager process (SANMAN). (1) A Guardian process with the name $ZZSMN that provides management access to the external ServerNet X and Y fabrics. (2) A Windows NT process that configures and maintains ServerNet switches within a Windows NT cluster.

fabric. A complex set of interconnections through which there can be multiple and (to the user) unknown paths from point to point. The term *fabric* is used to refer to the X or Y portion of the ServerNet communications network; for example, the X fabric.

factory-installed operating system. The version of the operating system image that HP creates having a CONFTEXT configuration file, OSIMAGE file, and configuration database that matches your order. Your system is shipped with this version installed in the system subvolume $SYSTEM.SYS00.

fan. A component within an HP NonStop™ S-series system enclosure that circulates air into the enclosure to help maintain optimal temperature. Each NonStop S-series system enclosure contains two fans.

far gateway. A short code sequence that accomplishes the transition to privileged mode for legitimate calls to callable procedures that are located in a different direct jump area. Typically, SCr (system code, RISC) is the target area. See also direct jump area, far jump, and gateway.

far jump. A sequence of RISC instructions that permits crossing the boundaries of the 256-megabyte direct jump areas in virtual memory. Such sequences are necessary, for example, when calling into system code from user code, because the two are located in different direct jump areas. The sequence ends with a JR (Jump via Register) RISC instruction.

fastLoad. An optimization that allows the loader to avoid reading symbols and binding symbolic references when loading a program or dynamic-link library in an environment equivalent to that of a previous load. See also preset, cached bindings, and library import characterization (LIC).

Fast Ethernet ServerNet adapter (FESA). A single-ported ServerNet adapter that supports 100-megabit/second (Mbps) or 10-Mbps Ethernet data transfer rates on an HP NonStop™ S-series server. The 3863 FESA installs directly into an available I/O port.

fault tolerance. The ability of a computer system to continue processing despite the failure of any single software or hardware component within the system.

FCSA. See Fibre Channel ServerNet adapter (FCSA).

FC switch. See Fibre Channel switch (FC switch).

FDC. See flexible disk configuration (FDC).

feature-test macro. In C and C++ programs, a symbol that, if defined in a program's source code, includes specific other symbols from a header within that program's source code and makes those symbols visible.

feeder circuit. The circuit conductors installed between the facility's main service entrance and the power distribution panels (PDPs) that supply the branch circuits.

ferrule. A cylindrical end terminal sometimes used on resistors, cartridge fuses, and other parts to permit quick insertion and removal from holders that have corresponding spring contacts.

FESA. See Fast Ethernet ServerNet adapter (FESA).

Fiber Optic Extension (FOX). Refers to two products, FOX II and ServerNet/FX, which allow you to create high-speed (up to 4 megabytes/second) networks of as many as 14 systems connected by dual fiber-optic cables.

fiber-optic plug-in card (F-PIC). A plug-in card (PIC) for the 6760 ServerNet device adapter (ServerNet/DA) that uses a fiber-optic interface to connect an HP NonStop™ S-series system to external disk drives and to some tape drives that contain a back-end board (BEB) that translates fiber-optic signals from the F-PIC into SCSI commands and information for the tape drive. See also plug-in card (PIC) and SCSI plug-in card (S-PIC).

fiber optics. A medium for data transmission that conveys light or images through very fine, flexible, glass or plastic fibers. Fiber-optic cables (light guides) are a direct replacement for conventional coaxial and wire pairs.

fiber-optic ServerNet addressable controller (F-SAC). A ServerNet addressable controller (SAC) that is contained within a fiber-optic plug-in card (F-PIC).

fiber-optic ServerNet cable. A ServerNet cable that uses fiber optics to transmit data. HP NonStop™ S-series servers support two types of fiber-optic ServerNet cables: multimode fiber-optic (MMF) ServerNet cable and single-mode fiber-optic (SMF) ServerNet cable.

Fibre Channel ServerNet adapter (FCSA). A ServerNet adapter that transmits data between an HP NonStop™ S-series server and Fibre Channel storage devices. This ServerNet adapter is installed in an IOAM enclosure.

Fibre Channel switch (FC switch). Networking hardware that can connect an Enterprise Storage System and an HP NonStop™ S-series server. This switch allows any-to-any connectivity.

field. In a structured programming language, an addressable entry within a data structure. The term *field* is sometimes used to mean member.

field-programmable gate array (FPGA). A programmable integrated circuit that can be customized to perform specific functions.

field-replaceable unit (FRU). A unit that can be replaced in the field only by qualified personnel trained by HP and cannot be replaced by customers. A unit is classified as a FRU because of safety hazards such as weight, size, sharp edges, or electrical potential; contractual agreements with suppliers; or national or international standards. See also customer-replaceable unit (CRU).

FIFO. A type of Open System Services (OSS) special file that is always read and written in a first-in, first-out manner.

FIFO special file. See FIFO.

file. An object to which data can be written or from which data can be read. A file has attributes such as access permissions and a file type. In the Open System Services (OSS) environment, file types include regular file, character special file, block special file, FIFO, and directory. In the Guardian environment, file types include disk files, processes, and subdevices.

file class. The property of an Open System Services (OSS) file indicating access permissions for a process related to the owner, group, or other identification of the process. See also file group class, file other class, and file owner class.

file description. See open file description.

file descriptor. In the Open System Services (OSS) file system, the nonnegative integer that uniquely identifies a single open of a file to a running process. Each file descriptor is associated with an open file description that contains data about the file.

file group class. The property of an Open System Services (OSS) file indicating access permissions for a process related to the group ID of the process. A process is in the file group class of a file if both:

- The process is not a member of the file owner class for the file.
- The process has an effective group ID or supplementary group ID that is the same as the group ID associated with the file.

file identifier. In the Guardian environment, the portion of a filename following the subvolume name. In the Open System Services (OSS) environment, a file identifier is a portion of the internal information used to identify a file in the OSS file system (an inode number). The two identifiers are not comparable.

file link count. The total number of directory entries for an Open System Services (OSS) file within an HP NonStop™ node.

file mode. For an Open System Services (OSS) process, a field in the `stat` structure for a specific file that describes the type and characteristics of the file and contains the access permission bits for the file.

file mode creation mask. A mask associated with an Open System Services (OSS) process and used when the process creates a file. Bits set in this mask are cleared in the access permission bits for the file.

filename. In the Open System Services (OSS) environment, a component of a pathname containing any valid characters other than slash (/) or null. See also file name.

file name. A string of characters that uniquely identifies a file.

In the PC environment, file names for disk files normally have at least two parts (the disk name and the file name); for example, B:MYFILE.

In the Guardian environment, disk file names include a node name, volume name, subvolume name, and file identifier; for example, \NODE.$DISK.SUBVOL.MYFILE.

In the Open System Services (OSS) environment, a file is identified by a pathname; for example, `/usr/john/workfile`. See also filename.

file other class. The property of an Open System Services (OSS) file indicating access permissions for a process related to the user ID and group ID of the process. A process is in the file other class of a file if both:

- The process is not a member of the file owner class for the file.
- The process is not a member of the file group class for the file.

file owner class. The property of an Open System Services (OSS) file indicating access permissions for a process related to the user ID of the process. A process is in the file owner class of a file if the process has an effective user ID that is the same as the user ID (owner) associated with the file.

file permission bits. Information about an Open System Services (OSS) file that is used, along with other information, to determine whether a process or user has read, write, or execute/search permission to that file. The bits are divided into three parts: owner, group, and other. Each part is used with the corresponding file class of processes.

file serial number. A number that uniquely identifies a file within its file system.

fileset. In the Open System Services (OSS) environment, a set of files with a common mount point within the file hierarchy. A fileset can be part or all of a single virtual file system.

On an HP NonStop™ system, the Guardian file system for a node has a mount point and is a subset of the OSS virtual file system. The entire Guardian file system therefore could be viewed as a single fileset. However, each volume and each process of subtype 30 within the Guardian file system is actually a separate fileset.

The term *file system* is often used interchangeably with *fileset* in UNIX publications.

file system. In the Open System Services (OSS) environment, a collection of files and file attributes. A file system provides the namespace for the file serial numbers that uniquely identify its files. Open System Services provides a file system (see also ISO/IEC IS 9945-1:1990 [ANSI/IEEE Std. 1003.1-1990], Clause 2.2.2.38); the Guardian application program interface (API) provides a file system; and OSS Network File System (NFS) provides a file system. (OSS NFS filenames and pathnames are governed by slightly different rules than OSS filenames and pathnames.) Within the OSS and OSS NFS file systems, filesets exist as manageable objects.

On an HP NonStop™ system, the Guardian file system for a node is a subset of the OSS virtual file system. Traditionally, the API for file access in the Guardian environment is referred to as the *Guardian file system*.

In some UNIX and NFS implementations, the term *file system* means the same thing as *fileset*. That is, a file system is a logical grouping of files that, except for the root of the file system, can be contained only by directories within the file system. See also fileset.

File Transfer, Access, and Management (FTAM). The Open Systems Interconnection (OSI) standard developed by the International Organization for Standardization (ISO) for network file exchange and management services.

file transfer protocol (FTP). (1) The Internet-standard, high-level protocol for transferring files from one machine to another. The server side requires the client to supply a logon identifier and password before it honors requests. FTP makes no assumptions about the file-naming structure of the source and destination systems, and it allows the file names of each system to be represented in the vernacular. (2) The application used to send complete files over Transmission Control Protocol/Internet Protocol (TCP/IP) services.

filler panel. A blank faceplate that is installed in place of a ServerNet adapter or plug-in card (PIC) to ensure proper ventilation.

fingerprint. A unique identifier calculated for a file and displayed in hexadecimal format.

FIPS. A Federal Information Processing Standard of the United States government.

FIPS 151-1. The Federal Information Processing Standard that specifies the requirements for conformance to an older draft of POSIX.1 (IEEE Std. 1003.1-1988) than the version adopted as ISO/IEC IS 9945-1:1990 and imposes some additional requirements.

FIPS 151-2. The Federal Information Processing Standard that specifies the requirements for conformance to POSIX.1 as ISO/IEC IS 9945-1:1990 and imposes some additional requirements.

FIR. See FRU information record (FIR).

FIRINIT. A diagnostic task used to update the communications line interface processor (CLIP) FRU information record (FIR) that is kept in the ServerNet wide area network (SWAN) concentrator CLIP flash memory.

FIRMUP. A diagnostic task used to update the copy of the Portable Silicon Operating System (pSOS) system product embedded kernel that is kept in the ServerNet wide area network (SWAN) concentrator communications line interface processor (CLIP) flash memory.

firmware. Code in memory that is necessary for the power-up initialization and communication with a host or device. The software for components of the ServerNet architecture (for example, an adapter) is called firmware. Some firmware for ServerNet components is downloaded when the system or component is loaded.

fixed process configuration. Using Subsystem Control Facility (SCF) to configure a generic process to always start in the first available processor (that is, to be fault tolerant).

flag. In a UNIX or Open System Services (OSS) command, a character sequence that begins with a hyphen and is processed as a unit.

flash memory. A type of memory that contains essential firmware and nonessential firmware.

flash PROM. A type of programmable read-only memory (PROM) that is electrically reprogrammable.

flat segment. A type of logical segment. Each flat segment has its own distinct address range within the process address space that never overlaps the range of any other allocated segments. Thus all allocated flat segments for a process are always available for use concurrently. See also logical segment and selectable segment.

flexible disk configuration (FDC). Configuration of an internal mirrored disk volume such that no single processor enclosure or I/O enclosure can result in loss of online access to the volume.

foreground process. An Open System Services (OSS) process that belongs to a foreground process group.

foreground process group. In the Open System Services (OSS) environment, a process group whose members have privileges for access to their controlling terminal that are denied to processes in background process groups of that terminal. Each session with a controlling terminal has only one foreground process group for that terminal. Contrast with background process group.

foreground process group ID. In the Open System Services (OSS) environment, the process group ID of a foreground process group.

four-lane link. The four single-mode fiber-optic (SMF) ServerNet cables that connect the two HP NonStop™ Cluster Switches on the same external fabric (for example, X1 and X2) in a split-star topology.

FOX. See Fiber Optic Extension (FOX).

FOXMON. See FOX monitor process.

FOX monitor process. The Fiber Optic Extension (FOX) monitor process for the ServerNet/FX adapter subsystem. The process name is $ZZFOX.

FOX ring. The fiber-optic cabling that connects the nodes in a Fiber Optic Extension (FOX) cluster. This term is also used to refer to the topology of a FOX network.

FPGA. See field-programmable gate array (FPGA).

F-PIC. See fiber-optic plug-in card (F-PIC).

frame. (1) An assembly of sheet-metal parts that is an integral part of an enclosure and might contain peripherals or a chassis, depending on the type of enclosure. The frame enables the enclosures to be stacked and has provisions for routing and securing cables. The frame of an enclosure has dimensions that conform to an industry-standard 19-inch rack. (2) A unit of transmission in some data communications protocols, usually containing header, data, and checksum fields. (3) In NonStop S-series processors, a 4096-byte unit of physical memory; also called a physical page.

frame base. An assembly consisting of casters, leveling pads, and frame sheet metal that is an integral part of a base enclosure.

free list. In the Open System Services (OSS) file system, the list of available inodes that can be allocated to files.

frequency. The number of complete cycles/second of sinusoidal variation. For alternating-current (AC) power lines, the most common frequencies are 60 hertz and 50 hertz.

FRU. See field-replaceable unit (FRU).

FRU information record (FIR). A collection of information that every field-replaceable unit (FRU) carries with it, such as part number, revision, track ID, and media access control (MAC) address.

F-SAC. See fiber-optic ServerNet addressable controller (F-SAC).

FTAM. See File Transfer, Access, and Management (FTAM).

FTP. See file transfer protocol (FTP).

gateway. (1) A device used to convert the message protocol of one network to that of another. (2) A short code sequence that accomplishes the transition to privileged mode for legitimate calls to callable procedures. See also far gateway.

GB. See gigabyte (GB).

GCSC. See Global Customer Support Center (GCSC).

general-purpose register (GPR). One of a small number of undedicated high-speed memory locations in a processor.

generic process. A process created and managed by the Kernel subsystem; also known as a system-managed process. A common characteristic of a generic process is persistence.

GESA. See Gigabit Ethernet ServerNet adapter (GESA).

G4SA. See Gigabit Ethernet 4-Port ServerNet adapter (G4SA).

Gigabit Ethernet ServerNet adapter (GESA). A single-port ServerNet adapter that provides 1000 megabits/second (Mbps) data transfer rates between HP NonStop™ S-series systems and Ethernet LANs. A GESA can be directly installed in slots 51 through 54 of an I/O enclosure and slots 53 and 54 of a processor enclosure.

Two versions of the GESA are available:

- 3865 GESA-C (T523572): a single-port copper version compliant with the 1000 Base-T standard (802.3ab)
- 3865 GESA-F (T523572): a single-port fiber version compliant with the 1000 Base-SX standard (802.z)

Gigabit Ethernet 4-Port ServerNet adapter (G4SA). A multiport ServerNet adapter that provides 1000 megabits/second (Mbps) data transfer rates between HP NonStop™ S-series systems and Ethernet LANs. The G4SA is installed in slots 1, 2, 3, 4, and 5 of an I/O adapter module (IOAM). There are two IOAMs in an I/O adapter module enclosure (IOAM enclosure), so a total of 10 G4SAs can be installed in an enclosure. The G4SA is the only LAN adapter supported for the IOAM enclosure. Although the G4SA supersedes the Ethernet 4 ServerNet adapter (E4SA), Fast Ethernet ServerNet adapter (FESA), and the Gigabit Ethernet ServerNet adapter (GESA), it cannot be installed on an HP NonStop™ S-series system.

gigabyte (GB). A unit of measurement equal to 1,073,741,824 bytes (1024 megabytes). See also kilobyte (KB), megabyte (MB), and terabyte (TB).

Global Customer Support Center (GCSC). A support organization that provides telephone and remote diagnostic support for HP customers. GCSCs are located all over the world. See also Online Support Center (OSC).

global offset table (GOT). A table of indirect addresses of data, including function descriptors, that might reside in a different loadfile.The GOT is an artifact of the native compiler.

globalized. The import-control characteristic of a loadfile that allows it to import symbols from any loadfile in the loadList of the program with which it is loaded. When those loadfiles offer multiple definitions of the same symbol, those loadfiles are searched in loadList sequence and the first definition found takes precedence. See also searchList.

globally unique ID (GUID). A unique, read-only number stored in nonvolatile memory (SEEPROM) on a ServerNet II Switch at the time of manufacture. The GUID also appears on the bar code label. This number can be used programmatically to identify the switch.

GOT. See global offset table (GOT).

GPR. See general-purpose register (GPR).

graphical user interface (GUI). A user interface that offers point-and-click access to program functions.

ground. A conducting connection, whether intentional or accidental, between an electrical circuit and either the earth or some conducting body that serves in place of the earth, such as an underground metal water pipe, structural steel, or a ground rod driven into the earth. See also earth ground.

grounded. Connected to earth or to some conducting body that serves in place of the earth.

grounded conductor. A system or circuit conductor that is intentionally grounded.

ground fault. Any undesired current path from a point of differing potential to ground.

ground fault interrupter. A device that interrupts the electric current to the load when a fault current to ground exceeds a predetermined value that is less than that required to operate the overcurrent protection device of the supply circuit.

grounding conductor. A conductor used to connect equipment of the grounded circuit of a wiring system to one or more earth-grounding electrodes. See also earth-grounding electrode.

group. (1) The set of all objects accessible by a pair of service processors (SPs) located in the processor multifunction (PMF) customer-replaceable unit (CRU). In an HP NonStop™ S-series server, a system enclosure has exactly one group. (2) In the Open System Services (OSS) environment, a set of user IDs with the same group ID.

group database. A database on a node that contains the group name, group ID, and user names for each group using that node.

group ID. The nonnegative integer used to identify a group of users of an HP NonStop™ network node. Each user of a node is a member of at least one group. When the identity of a group is associated with an Open System Services (OSS) process, a group ID value is referred to as one of:

- Real group ID
- Effective group ID
- Supplementary group ID
- Saved-set group ID

group list. An Open System Services (OSS) process attribute that is used with the effective group ID of the process to determine the file access permissions for the process.

GRT. See Guided Replacement Toolkit (GRT).

Guardian. An environment available for interactive or programmatic use with the HP NonStop™ Kernel operating system. Processes that run in the Guardian environment usually use the Guardian system procedure calls as their application program interface. Interactive users of the Guardian environment usually use the HP Tandem Advanced Command Language (TACL) or another HP product's command interpreter. Contrast with Open System Services (OSS).

Guardian environment. The Guardian application program interface (API), tools, and utilities.

Guardian services. An application program interface (API) to the HP NonStop™ Kernel operating system, plus the tools and utilities associated with that API. This term is synonymous with *Guardian environment*. See also Guardian.

Guardian user ID. See HP NonStop™ Kernel user ID.

GUI. See graphical user interface (GUI).

GUID. See globally unique ID (GUID).

guided procedure. A software tool that assists you in performing complex configuration or replacement tasks on an HP NonStop™ S-series server. OSM guided procedures are launched through actions in the OSM Service Connection. TSM guided procedures are accessible from the Start menu on your system console. Guided procedures include Replace PMF, IOMF, SNDA, SEB, or MSEB. Some names vary between OSM and TSM. For example, Add Node to ServerNet Cluster is the OSM equivalent of Configure ServerNet Node in TSM.

Guided Replacement Toolkit (GRT). A software product that guides you through online replacement of these customer-replaceable units (CRUs) on HP NonStop™ S-series systems: I/O multifunction (IOMF) CRUs, power supplies, processor multifunction (PMF) CRUs, and 6760 ServerNet device adapters. GRT is used only with older versions of TSM server software. If you are replacing a CRU in a system running TSM server version T7945AAX (shipped with G06.13) or later, use the appropriate guided procedure.

hard link. In the Open System Services (OSS) file system, the relationship between two directory entries for the same file. A hard link acts as an additional pointer to a file. A hard link cannot be used to point to a file in another fileset. Contrast with symbolic link.

hard reset. An action performed on an HP NonStop™ Cluster Switch (model 6770) and HP NonStop™ ServerNet Switch (model 6780) that reinitializes the router-2 ASIC within the switch, disrupting the routing of ServerNet messages through the switch for several minutes. When the hard reset is finished, the paths are restored automatically.

harmonic. The sinusoidal component of an alternating-current (AC) voltage that is a multiple of the waveform frequency.

harmonic distortion. Harmonics that change an alternating-current (AC) waveform from sinusoidal to complex.

header. An object that, when specified for inclusion in a program's source code, causes the program to behave as if the statement including the header were actually a specific set of other programming statements. A header contains coded information that provides details (such as data item length) about the data that the header precedes.

In an Open System Services (OSS) program, a header is the name of a file known to the run-time library used by a process. In a Guardian environment C language program, a header is the file identifier for a file known to the run-time library used by a process.

hertz (Hz). A unit of frequency. One hertz equals one cycle/second.

high frequency. A Federal Communications Commission (FCC) designation for a frequency in the range 30 through 300 megahertz, corresponding to a decametric wave in the range 10 through 100 meters.

high PIN. A process identification number (PIN) that is greater than 255. Contrast with low PIN.

hop count. The number of routers that form a route between a ServerNet source and ServerNet destination. Hop count is used to determine the best route. If two alternate routes have the same time factor, the path with the lower hop count is the better route.

host database. An SQL database maintained for the host system and containing information about requests, software inputs, snapshots, targets, and profiles.

host system. (1) A computer system that supports very large databases and does batch processing, usually for an entire network of smaller systems. (2) The central site on which the Distributed Systems Management/Software Configuration Manager (DSM/SCM) is managed, the Archive is maintained, and configuration revisions are built. The host system is also a target system.

HP NonStop™ Cluster Switch (model 6770). An assembly that routes ServerNet messages across an external fabric of a ServerNet cluster. The cluster switch consists of a ServerNet II Switch, an uninterruptible power supply (UPS), and AC transfer switch, and it can be packaged in a switch enclosure or in a 19-inch rack. The cluster switch is used with star, split-star, and tri-star topologies. See also HP NonStop™ ServerNet Switch (model 6780).

HP NonStop™ Kernel Open System Services (OSS). The product name for the OSS environment. See also Open System Services (OSS).

HP NonStop™ Kernel operating system. The operating system for HP NonStop systems.

HP NonStop™ Kernel user ID. A user ID within an HP NonStop system. The Guardian environment normally uses the structured view of this user ID, which consists of either the *group-number, user-number* pair of values or the *group-name.user-name* pair of values. For example, the structured view of the super ID is (255, 255). The Open System Services (OSS) environment normally uses the scalar view of this user ID, also known as the UID, which is the value (*group-number* * 256) + *user-number*. For example, the scalar view of the super ID is (255 * 256) + 255 = 65535.

HP NonStop™ K-series servers. The set of servers in the HP NonStop servers having product numbers beginning with the letter *K*. These servers run the HP NonStop Kernel operating system, but they do not implement the ServerNet architecture.

HP NonStop™ Open System Management (OSM) Interface. Replacement for TSM as the system management tool of choice for NonStop S-series servers. OSM provides the same functionality as TSM while overcoming limitations of TSM. OSM is required for support of new functionality released in G06.21 and later.

HP NonStop™ S700 server. A special configuration of HP NonStop S-series server that is limited to one processor enclosure and a maximum of two I/O enclosures. A matched pair of any model of PMF CRU can be used in a NonStop S700 server.

HP NonStop™ S7000 server. The first server in a product line of servers (HP NonStop S-series servers) that implement the ServerNet architecture and run the HP NonStop Kernel operating system.

HP NonStop™ S7400 server. A model of HP NonStop S-series server that provides a midrange upgrade option for migrating from an HP NonStop K-series server or a NonStop S7000 server.

HP NonStop™ S7600 server. A model of HP NonStop S-series server that implements the ServerNet architecture and runs the HP NonStop Kernel operating system. The NonStop S7600 PMF CRU is based on the NonStop S74000 PMF CRU, and the NonStop S7600 server supports all S-series hardware products that are compatible with the NonStop S74000 server.

HP NonStop™ S7800 server. A model of the HP NonStop S-series server that provides a midrange upgrade option from a NonStop S7600 server or earlier midrange NonStop S-series server.

HP NonStop™ S70000 server. See HP NonStop™ Sxx000 server.

HP NonStop™ S72000 server. See HP NonStop™ Sxx000 server.

HP NonStop™ S74000 server. See HP NonStop™ Sxx000 server.

HP NonStop™ S76000 server. See HP NonStop™ Sxx000 server.

HP NonStop™ S86000 server. See HP NonStop™ Sxx000 server.

HP NonStop™ Sxx000 server. Any server in a family of high-performance servers (HP NonStop S-series servers) that implement the ServerNet architecture and run the HP NonStop Kernel operating system. This family includes the NonStop S70000, S72000, S74000, S76000, and S86000 servers.

HP NonStop™ ServerNet Cluster (ServerNet Cluster). The product name for the collection of hardware and software components that constitute a ServerNet cluster.

HP NonStop™ ServerNet Switch (model 6780). The cluster switch used in the layered topology. The 6780 switch consists of a switch logic board, a midplane, plug-in cards, power supplies, and fans. See also HP NonStop™ Cluster Switch (model 6770).

HP NonStop™ S-series servers. The HP NonStop servers having product numbers beginning with the letter *S*. These servers implement the ServerNet architecture and run the HP NonStop Kernel operating system.

HP NonStop™ servers. The entire line of HP NonStop servers, including NonStop K-series and NonStop S-series servers.

HP NonStop™ Storage Management Foundation (SMF). A subsystem used by the storage subsystem that facilitates automation of storage management tasks by providing location-independent naming, storage pools, and virtual disks on HP NonStop™ S-series systems.

HP NonStop™ System RISC Model D processor (NSR-D processor). The model designation for the TNS/R processor used in the HP NonStop S7400 server.

HP NonStop™ System RISC Model E processor (NSR-E processor). The model designation for the TNS/R processor used in the HP NonStop S7600 server.

HP NonStop™ System RISC Model G processor (NSR-G processor). The model designation for the TNS/R processor used in the HP NonStop S70000 server.

HP NonStop™ System RISC Model T processor (NSR-T processor). The model designation for the TNS/R processor used in the HP NonStop S72000 server.

HP NonStop™ System RISC Model V processor (NSR-V processor). The model designation for the TNS/R processor used in the HP NonStop S74000 server.

HP NonStop™ System RISC Model W processor (NSR-W processor). The model designation for the TNS/R processor used in the HP NonStop S7000 server.

HP NonStop™ System RISC Model X processor (NSR-X processor). The model designation for the TNS/R processor used in the HP NonStop S76000 server.

HP NonStop™ System RISC Model Y processor (NSR-Y processor). The model designation for the TNS/R processor used in the HP NonStop S86000 server.

HP NonStop™ TCP/IP. The HP implementation of Transmission Control Protocol/Internet Protocol (TCP/IP) for the HP NonStop servers. See also Parallel Library TCP/IP.

HP NonStop™ TCP/IP process. An HP product that supports the Transmission Control Protocol/Internet Protocol (TCP/IP) layers. TCP/IP processes are used together with the communications line interface processor (CLIP) pNA+ to provide the transport layer between wide area network (WAN) I/O processes and data link control (DLC) tasks, between ConMgr and the Simple Network Management Protocol (SNMP) task, between the WANBoot process and BOOTP tasks, and between an OSM or TSM process and a DIAG task.

HP NonStop™ TCP/IP subsystem. A subsystem that allows the use of HP NonStop TCP/IP to access an HP NonStop S-series host from Macintosh computers, personal computers, and UNIX workstations. Applications running on a NonStop S-series system or in an Expand network can transparently exchange data with NonStop TCP/IP devices.

HP NonStop™ TCP/IPv6. An HP product that adds IP version 6 (IPv6) functionality to the parallel library TCP/IP product. IPv6 is a TCP/IP protocol that extends the IP version 4 (IPv4) of 32 bits to 128 bits. NonStop TCP/IPv6 can be run in three modes: INET (only IPv4 and is a direct replacement for parallel library TCP/IP), INET 6 (only IPv6), and Dual (both IPv4 and IPv6 communications).

HP NonStop™ Technical Library (NTL). The application for accessing, searching, and viewing technical publications and support information for the HP NonStop server. NTL replaces Total Information Manager (TIM).

HP NonStop™ Transaction Management Facility (TMF). HP software that provides transaction protection and database consistency in demanding online transaction processing (OLTP) and decision-support environments. It gives full protection to transactions that access distributed SQL and Enscribe databases, as well as recovery capabilities for transactions, online disk volumes, and entire databases

HP Tandem Advanced Command Language (TACL). The user interface to the HP NonStop™ Kernel operating system. The TACL product is both a command interpreter and a command language. Users can write TACL programs that perform complex tasks or provide a consistent user interface across independently programmed applications.

HP Tandem Failure Data System (TFDS). A diagnostic tool that is a component of the HP NonStop™ Kernel operating system. The TFDS tool isolates software problems and provides automatic processor-failure data collection, diagnosis, and recovery services.

HP Transaction Application Language (TAL). A systems programming language with many features specific to stack-oriented TNS systems.

hybrid shared run-time library (hybrid SRL). A shared run-time library (SRL) that has been augmented by the addition of a dynamic section that exports the SRL's symbols in a form that can be used by position independent code (PIC) clients. A hybrid SRL looks like a dynamic-link library (DLL) to PIC clients (except it cannot be loaded at other addresses and cannot itself link to DLLs). The code and data in the SRL are no different in a hybrid SRL, and its semantics for non-PIC clients are unchanged.

Hz. See hertz (Hz).

I18N. See internationalization.

IBC. See in-band control (IBC).

ICMP. See Internet control message protocol (ICMP).

identifier. A unique name; for example, TANDEM^FILES^TO^COPY; in the CONFTEXT file that refers to a text string (one or more file names as given in the CONFAUX file). When Distributed Systems Management/Software Configuration Manager (DSM/SCM) encounters an identifier, it substitutes the text string for the identifier.

ideogram. See ideograph.

ideograph. A character or symbol representing a word or idea. Some writing systems, such as Japanese and Chinese, use thousands of ideographs. An ideograph is sometimes called an *ideogram*.

IEC. International Electrotechnical Committee. IEC is a professional organization that creates or adopts standards for computer hardware, environments, and physical interconnections.

IEEE. Institute of Electrical and Electronics Engineers. IEEE is a professional organization whose committees develop and propose computer standards that define the physical and data link protocols of entities such as communication networks.

IEEE 802.3 protocol. Institute of Electrical and Electronics Engineers (IEEE) standard defining the hardware layer and transport layer of (a variant of) Ethernet. The maximum segment length is 500 meters, and the maximum total length is 2.5 kilometers. The maximum number of hosts is 1024. The maximum packet size is 1518 bytes.

impedance. The total opposition (that is, resistance and reactance) a circuit provides to the flow of alternating current at a given frequency.

implementation-defined. Not specified by a standard. A correct value or behavior that is implementation-defined can vary from system to system and therefore might represent a feature or facility that cannot be ported.

implicit library. A library supplied by HP that is available in the read-only and execute-only globally mapped address space shared by all processes without being specified to the linker or loader. See also TNS system library and public library.

implicit library import library (imp-imp). See import library.

implied user library. A method of binding TNS object files that have more than 16 code segments. Segments 16 through 31 are located in the user code (UC) space but are executed as if they were segments 0 through 15 of the user library (UL) code space. This method precludes the use of a user library. Binder now supports 32 segments of UC space concurrently with 32 segments of UL code space, so the implied user library method is not needed in new or changed TNS applications.

import. To refer to a symbol definition from another loadfile. A loadfile imports a symbol definition when it needs a data item or function having that symbolic name.

import control. The characteristic of a loadfile that determines from which other loadfiles it can import symbol definitions. The programmer sets a loadfile's import control at link time. That import control can be localized, globalized, or semiglobalized. A loadfile's import control governs the way the linker and loader construct that loadfile's searchList and affects the search only for symbols required by that loadfile.

import library. A file that represents a dynamic-link library (DLL) and can substitute for it as input to the linker. Import libraries facilitate linking on auxiliary platforms (that is, PCs) where it is inconvenient to store the actual DLLs.

in-band control (IBC). A symbol-based communications protocol for communicating management information across a ServerNet link without interfering with any application traffic in the network. In ServerNet II architecture, IBC traffic uses standard ServerNet packets. ServerNet I architecture uses the Illegal Symbol variation of IBC, which uses a subset of the available symbols to convey control information from one node to another. The symbol subset chosen is from the group of symbols that are not used for passing data. These symbols are usually considered illegal or unused.

incident report. A report sent by the OSM or TSM server software to the respective OSM or TSM Notification Director. If remote notification (dial-out) is configured, the Notification Director forwards incident reports to a service provider. The three types of incident reports are problem incident reports, periodic incident reports, and software configuration incident reports.

incremental discovery. Discovery of an HP NonStop™ S-series server when the OSM or TSM client software has locally saved information but where configuration changes have been made on the server since that information was saved.

indicator lights. Two light-emitting diodes (LEDs) on a customer-replaceable unit (CRU) that indicate the status of the unit. The red or amber indicator light is lit when the unit is not working properly. During startup, this light can indicate that the unit is not yet functioning. The green indicator light is lit when the unit has proper power applied. See also light-emitting diode (LED).

inductive reactance. Resistance at a frequency that is caused by the inductance of a coil or circuit.

initial discovery. Discovery of an HP NonStop™ S-series server with which the OSM or TSM client software has had no prior contact and for which the client software has no locally saved information.

initialization. The process of defining a new Distributed Systems Management/Software Configuration Manager (DSM/SCM) target system, including giving it a name, setting up default values used when processing requests, and creating the first software revision (list of products) for the system.

initial software revision. The software revision on a target system when it is first brought into the Distributed Systems Management/Software Configuration Manager (DSM/SCM) environment. The DSM/SCM host database must be initialized with information about the initial software revision. The initial software revision is then used as a baseline upon which new software revisions are based.

inode. A data structure that stores the location of an Open System Services (OSS) file.

inode number. A unique identifier within the Open System Services (OSS) file system of an instance of an OSS file. The inode number identifies the instance within the file system catalogs.

input/output process (IOP). A running program (part of the HP NonStop™ Kernel operating system) that manages the I/O functions for one or more ServerNet addressable controllers (SACs) of the same type.

input source. The resource from which Subsystem Control Facility (SCF) accepts command input. SCF can accept input from a terminal or a disk file. The initial input source is determined by the form of the RUN command used to initiate SCF. At any time during an SCF session, the input source can be temporarily changed to execute a series of commands from a command file.

inrush current. The initial surge current demand of a load.

Inspect region. The region of a TNS object file that contains symbol tables for all blocks compiled with the SYMBOLS directive. The Inspect region is sometimes called the symbols region.

INSPSNAP. The program that provides a process snapshot file for the Inspect subsystem.

installation subvolume (ISV). A subvolume containing files that perform a specific function during the installation process, such as organizing documentation in a specific location, providing the components of the HP NonStop™ Kernel operating system image (OSIMAGE), and containing files that are used after the installation process.

installer. The person who installs the system equipment for a new system. This person also installs new equipment when additions are made to the system. This person can install software and perform system verification procedures as directed by the system planner, configuration planner, or support planner.

instruction processing unit (IPU). A processing unit that executes programs by fetching instructions from memory and executing them.

insulated ground. A grounding conductor with a dielectric (low-conductance) insulator around it to prevent inadvertent contact with metal conduits.

intelligent SCSI processor (ISP). The ServerNet addressable controller (SAC) that controls the small computer system interface (SCSI) bus.

interactive mode. A mode of operation that is characterized by having the same input and output device (a terminal or a process) for the session. If a terminal is used, a person enters a command and presses Return. If a process is used, the system interface waits for the process to send a request and treats the process in the same manner as a terminal. Contrast with noninteractive mode.

internal routing. The routing of packets within an HP NonStop™ S-series server. See also external routing.

internal ServerNet fabrics. The fabrics that link ServerNet devices within an HP NonStop™ S-series server. See also external ServerNet fabrics.

internal ServerNet X or Y fabric. The X or Y fabric that links ServerNet devices within an HP NonStop™ S-series server. See also fabric.

internationalization. The process of designing and coding software so that it can be adapted to meet the needs of different languages, cultures, and character sets, with the ability to handle various linguistic and cultural conventions. Internationalization methods enable the processing of character-based data independently of the underlying character encoding, allowing choice among character sets. Sometimes referred to as *I18N*, derived from the 18 letters between the initial *I* and the final *N* of the word *internationalization*. See also character set.

Internet address. The 32-bit address assigned to hosts that want to participate in the Internet using Transmission Control Protocol/Internet Protocol (TCP/IP). Internet addresses are an abstraction of physical hardware addresses, just as the Internet is an abstraction of physical networks. As assigned to the interconnection of a host to a physical network, an Internet address consists of a network portion and a host portion. See also IP address.

Internet control message protocol (ICMP). A maintenance protocol in the Transmission Control Protocol/Internet Protocol (TCP/IP) suite that is required in every TCP/IP implementation. The ICMP allows two nodes on an IP network to share IP status and error information. The ICMP is used by the ping utility to determine the readability of a remote system. See also IP address and ping.

Internet protocol (IP). A data communications protocol that handles the routing of data through a network, which typically consists of many different subnetworks. IP is connectionless. It routes data from a source address to a destination address. See also IP address.

interoperability. (1) Within an HP NonStop™ node, the ability to use the features or facilities of one environment from another. For example, the `gtacl` command in the Open System Services (OSS) environment allows an interactive user to start and use a Guardian tool in the Guardian environment. (2) Among systems from multiple vendors or with multiple versions of operating systems from the same vendor, the ability to exchange status, files, and other information. Product externals and end-user publications for the NonStop range of servers often use the term *connectivity* in this context. See also connectivity.

interprocessor communications (IPC). The exchange of messages between processors.

intrinsic library. See Shared Millicode Library.

instance. A particular case of a class of items or objects. For example, a process is defined as one instance of the execution of a program. Multiple processes might be executing the same program simultaneously. Also, instance data refers to global data of a program or library. Each process has its own instance of the data.

instance data. For each process using a dynamic-link library, a data segment area containing the global variables used by the library.

I/O adapter module (IOAM). A collection of modular components that provides I/O connectivity and can include ServerNet switch boards, Fibre Channel ServerNet adapters (FCSAs), fans, and power supplies. In the IOAM, each module is a logical entity that represents a single service domain.

I/O adapter module enclosure (IOAM enclosure). The sheet-metal carrier that is installed in a standard 19-inch rack and contains the IOAM components.

IOAM. See I/O adapter module (IOAM).

IOAM enclosure. See I/O adapter module enclosure (IOAM enclosure).

I/O cabinet. See I/O enclosure.

I/O controller. The hardware logic that controls computer I/O operations for a particular set of devices, such as disks, tapes, terminals, or communications lines. See also ServerNet addressable controller (SAC).

I/O enclosure. An HP NonStop™ S-series system enclosure containing one module, which includes ServerNet adapters, disk drives, components related to the ServerNet fabrics, and components related to electrical power and cooling for the enclosure. An I/O enclosure is identical to a processor enclosure except that it contains I/O multifunction (IOMF) customer-replaceable units (CRUs) instead of processor multifunction (PMF) CRUs.

IOMF CRU. See I/O multifunction (IOMF) CRU.

IOMF 2 CRU. See I/O multifunction (IOMF) 2 CRU.

I/O multifunction (IOMF) CRU. (1) An HP NonStop™ S-series customer-replaceable unit (CRU) that connects an I/O enclosure to a processor enclosure through a ServerNet cable and supplies power to the components within the IOMF CRU as well as redundantly to the disk drives, SCSI terminators, and ServerNet adapters in that enclosure. The IOMF CRU contains a power supply, a service processor (SP), a ServerNet router, an Ethernet controller, an external ServerNet port, and three SCSI ServerNet addressable controllers (S-SACs) in a single unit. (2) A collective term for both IOMF CRUs and IOMF 2 CRUs when a distinction between the two types of CRUs is not required.

I/O multifunction (IOMF) 2 CRU. An HP NonStop™ S-series customer-replaceable unit (CRU) that connects an I/O enclosure to a processor enclosure through a ServerNet cable and supplies power to the components within the IOMF 2 CRU as well as redundantly to the disk drives, SCSI terminators, and ServerNet adapters in that enclosure. The IOMF 2 CRU contains a power supply, a service processor (SP), a ServerNet router 2, an Ethernet controller, three configurable ServerNet ports, and three SCSI ServerNet addressable controllers (S-SACs) in a single unit. IOMF 2 CRUs are supported on G06.10 and later release version updates (RVUs).

IOP. See input/output process (IOP).

IP. See Internet protocol (IP).

IP address. An address that uniquely identifies a specific host system within a network to the Internet protocol (IP). An IP address consists of two parts: a network address, which identifies the network, and a local address, which identifies the host within the network. IP routes data between source and destination IP addresses.

iPAQ Desktop. An Internet-based computing model and business PC used for Internet access, mainstream computing, and running the Windows 2000 Professional operating system. The iPAQ Desktop is designed to reduce hardware and software conflicts. It eliminates Industry Standard Architecture (ISA)/Peripheral Component Interconnect (PCI) slots and uses Universal Serial Bus (USB) ports with the Windows 2000 or Windows ME operating system.

IPC. See interprocessor communications (IPC).

IPU. See instruction processing unit (IPU).

ISO. International Organization for Standardization. ISO is an international body that drafts, discusses, proposes, and specifies standards for network protocols. ISO is best known for its seven-layer reference model that describes the conceptual organization of protocols.

ISO is sometimes called the *International Standards Organization*. Although ISO is the official abbreviation, it does not correspond to the organization's name in any language.

ISO 646. An ISO standard for representing characters in languages based on the Roman alphabet. Like ASCII, ISO 646 uses only 7 bits of each 8-bit byte to represent data. Contrast with ISO 8859.

ISO 8859. A series of ISO standard 8-bit code sets used to represent languages based on many alphabets, including Roman, Greek, Cyrillic, Hebrew, Turkish, and Arabic. The ISO 8859 code sets are used in international applications that must be data transparent. ASCII is a subset of each of the ISO 8859 code sets.

ISO 10646. A universal coded character set that represents all characters and symbols from all commonly used scripts and languages.

ISO/IEC-conforming POSIX.1 application. An application that both:

- Uses only the facilities described in ISO/IEC IS 9945-1:1990 and approved conforming language bindings for any ISO or IEC standard
- Is documented as using only those facilities and approved conforming language bindings

isolated ground. A grounding conductor that directly connects the equipment ground through an isolated ground-type receptacle with the power system grounding point without any intermediate grounding points.

isolation transformer. A transformer containing electrostatic shields between the primary and secondary windings, with no direct electrical path between the primary and secondary windings.

ISP. See intelligent SCSI processor (ISP).

ISV. See installation subvolume (ISV).

JDS box. See ServerNet extender module (SEM).

job control. The Open System Services (OSS) features that allow processes to be stopped, continued, and moved from or to the background.

KB. See kilobyte (KB).

Kernel subsystem. In G-series release version updates (RVUs), the subsystem for configuration and management of the Subsystem Control Facility (SCF) subsystem managers that are generic processes, some system attributes, and the ServerNet X and Y fabrics.

Kernel subsystem manager process. The generic process that starts and manages other generic processes, some system attributes, and the ServerNet X and Y fabrics in G-series release version updates (RVUs). The $ZZKRN Kernel subsystem manager process is started and managed by the $ZPM persistence manager process.

kilobyte (KB). A unit of measurement equal to 1024 bytes. See also gigabyte (GB), megabyte (MB), and terabyte (TB).

K-series servers. See HP NonStop™ K-series servers.

L10N. See localization.

labeled dump. A token-by-token display-text representation of the Subsystem Programmatic Interface (SPI) command buffer or response buffer, as produced by the Subsystem Control Facility (SCF) commands DETAIL CMDBUFFER and DETAIL RSPBUFFER. The display text includes a labeled value for each token.

LAN. See local area network (LAN).

LANMAN. See LAN manager (LANMAN) process.

LAN manager (LANMAN) process. The process provided as part of the ServerNet local area network (LAN) Systems Access (SLSA) subsystem that starts and manages the SLSA subsystem objects and the LAN monitor (LANMON) process and assigns ownership of Ethernet adapters to the LANMON processes in the system. Subsystem Control Facility (SCF) commands are directed to the LANMON processes for configuring and managing the SLSA subsystem and the Ethernet adapters.

LANMON. See LAN monitor (LANMON) process.

LAN monitor (LANMON) process. The process provided as part of the ServerNet local area network (LAN) Systems Access (SLSA) subsystem that has ownership of the Ethernet adapters controlled by the SLSA subsystem.

late binding. At load time, binding a symbolic reference in a dynamic-link library (DLL) to a definition in a loadfile that appears on the program's loadList rather than the one found on the DLL's linker searchList. Late binding occurs in either of these cases:

- The loader resolves a symbol that is unresolved by any loadfile on the linker searchList.

- The loader binds a symbol in a DLL to the first definition it finds on the program's loadList, and this is not the first definition that was encountered on the linker searchList.

For localized loadfiles, the linker and loader searchLists are the same, so late binding does not occur.

layer number. See cluster switch layer number.

layered topology. The network topology for ServerNet clusters using the HP NonStop™ ServerNet Switch (model 6780). The layered topology can scale by adding cluster switch layers or zones. The layered topology supports up to four layers and three zones. See also star topology, split-star topology, and tri-star topology.

LB. See logic board (LB).

LC-LC cable. A multimode fiber cable that connects an Fibre Channel ServerNet adapter (FCSA) to a Fibre Channel storage device or a Fibre Channel switch (FC switch).

LC-SC cable. A multimode fiber cable that connects an modular ServerNet expansion board (MSEB) to an I/O adapter module (IOAM) or a Fibre Channel switch (FC switch).

ldev. See logical device number (ldev).

LED. See light-emitting diode (LED).

legacy system. An operating system that is not open but from which applications must be ported or users transferred.

libList. The list of libraries to be loaded along with a loadfile. When linking the loadfile, the linker constructs the libList from the names of libraries specified in the linker's command stream. It stores the libList within the loadfile.

library. A generic term for a collection of routines useful in many programs. An object code library can take the form of a linkfile to be physically included into client programs, it can be an OSS archive file containing several linkable modules, it can be a loadfile, or it can be a system-managed collection of preloaded routines. Source-code libraries fall outside the scope of this glossary. See also dynamic-link library (DLL) and shared run-time library (SRL).

library client. A program or another library that uses routines or variables from that library.

library file. See library.

library import characterization (LIC). A list of the export digests and relocation offsets of all the libraries used to resolve symbols in a loadfile. It allows the loader and operating system to determine when a file is being loaded in an environment equivalent to that found by the linker or to a previous load. (See fastLoad.) A LIC is generated and stored in the loadfile by the linker when a file is preset; it can be used in a subsequent load step to determine whether the loadfile's existing bindings can be reused. The operating system can also retain the bindings as modified when a loadfile is loaded and associate a LIC with those cached bindings, so that they can be reused when the same file is again loaded in an equivalent environment. (See cached bindings.)

LIC. See library import characterization (LIC).

LIF. See logical interface (LIF).

light-emitting diode (LED). A semiconductor device that emits light from its surface. Indicator lights are composed of LEDs. See also indicator lights.

line. The specific hardware path over which data is transmitted or received. A line can also have a process name associated with it that identifies an input/output process (IOP) or logical device associated with that specific hardware path.

linear load. Electrical loads for which the impedance is constant regardless of the voltage, so that if the voltage is sinusoidal, the current drawn is also sinusoidal.

line-handler process. See Expand line-handler process or Expand-over-ServerNet line-handler process.

line interface unit (LIU). A dual-ported unit consisting of two parts: a communications line interface processor (CLIP) and a line interface module (LIM). An LIU can communicate with either the primary or the backup processor, providing fault tolerance. When it is a component of the communications subsystem, an LIU communicates with either processor through either of a pair of communications interface units (CIUs).

link. In the Open System Services (OSS) file system, a directory entry for a file.

link count. In the Open System Services (OSS) file system, the number of directory entries that refer to a particular file.

linker. (1) The process or server that invokes the message system to deliver a message to some other process or server. (2) A programming utility, which combines one or more compilation units' linkfiles to create an executable loadfile for a native program or library.

linkfile. (1) For native C/C++ compilers in the Guardian environment, a command file for input to the `nld` or `ld` utility. (2) A file containing object code that is not yet ready to load and execute. Linkfiles are combined by means of a linker or binder to make an executable loadfile for a program or library. Compiling creates one linkfile per independent source module. Contrast with loadfile.

linking. The operation of collecting, connecting, and relocating code and data blocks from one or more separately compiled object files to produce a target object file.

link name. In the Open System Services (OSS) environment, the filename associated with a specific file within a directory. The length of a filename, and therefore the length of a link name, depends on the file system.

Linux. Linus Torvald's version of the UNIX operating system. See also `http://www.linux.org`.

listener. The process or server that is notified by the message system that a message from some other process or server is being delivered.

LIU. See line interface unit (LIU).

LMU. See logical memory unit (LMU).

load. (1) To transfer the HP NonStop™ Kernel operating system image or a program from disk into a computer's memory so that the operating system or program can run. (2) To insert a tape into a tape drive, which prepares it for a tape operation (read or write).

loadable library. A loadfile that offers functions and data to other loadfiles. In this manual, dynamic-link libraries and hybrid shared run-time libraries are libraries. A library cannot normally be invoked externally; for example, by a RUN command. Instead, it is invoked by calls or data references from client loadfiles.

loader. A programming utility that transfers a program into memory so it can run. The mechanism that brings loadfiles into memory for execution, maps them into virtual address space, and resolves symbol references among them. Synonyms include run-time loader and run-time linker. The loader for TNS and for TNS/R native programs and libraries that are not position-independent code (PIC) is part of the operating system. For PIC loadfiles, a loader called RLD works with the operating system to load programs and libraries.

loadfile. An executable object code file that is ready for loading into memory and executing on the computer. Loadfiles are further classified as executable programs (containing a main routine at which to begin execution of that program) or executable libraries (supplying routines or variables to multiple programs or separately loaded libraries). A TNS code file might be both a loadfile and a linkfile. Native code files are never both. Contrast with linkfile.

loadList. A list of all the libraries that must be loaded for a given loadfile to execute. A loadfile's loadList includes all the libraries in the given loadfile's libList plus all the libraries in those loadfiles' libLists, and so forth. It does not include the implicit libraries. The loadList order is the sequence in which these loadfiles are to be loaded when they are not already loaded by a previous operation. The loadList of the program includes all the loadfiles present in the process, in the order they were loaded.

local area network (LAN). A network that is located in a small geographical area and whose communications technology provides a high-bandwidth, low-cost medium to which low-cost nodes can be connected. One or more LANs can be connected to the system such that the LAN users can access the system as if their workstations were connected directly to it. Contrast with wide area network (WAN).

locale. In localization, the definition of the subset of a user's environment that depends on language and cultural conventions.

localization. The process of adapting computer interfaces, data, and documentation to the culturally accepted way of presenting information in the culture. Sometimes referred to as *L10N*, derived from the 10 letters between the initial *L* and the final *N* of the word *localization*.

localized. The import-control characteristic of a loadfile that allows it to import symbols only from the loadfile itself followed by the libraries in its libList, libraries that those libraries reexport, and from these, any successions of reexported libraries. See also register-exact point and searchList.

local mount. In the Network File System (NFS), a mount that attaches the fileset associated with a server to the specified mount point within the local directory hierarchy. The local mount is visible within the NFS subsystem and makes the files associated with the server available through the path associated with the local mount point.

local node. See local system.

local operator. The person who performs routine system operations, such as starting and stopping the system, loading and unloading tapes, and changing the air filter. The local operator is normally the operator of the asynchronous system console for the node. See also operator.

local processor. A processor in the same node as the ServerNet cluster monitor process (SNETMON) that is reporting status about the processor.

local system. (1) An on-site system or a system that is geographically near the user or operator. (2) From the perspective of a particular SNETMON, the system or node on which that SNETMON is running. (3) From the perspective of a system console operator, the system to which the operator is logged on. Contrast with remote system.

logical device name. The name assigned to an I/O process during its configuration. Other processes use the logical device name when issuing Guardian procedure calls to the I/O process.

logical device number (ldev). For a device or named process, the index into the destination control table (DCT); or more precisely, into the named resource list (NRL). Each device or named process in the system has an ldev.

logical disk volume. A hardware device or device pair that provides persistent, highly accessible storage for data on a medium that is either magnetic or optical.

logical interface (LIF). A process that allows an application or another process to communicate with data communications hardware.

logical memory unit (LMU). A group of four memory units. Memory on a processor and memory board (PMB) is divided into two LMUs. One LMU contains memory units in slots MS1 through MS4. The other LMU contains memory units in slots MS5 through

MS8. An LMU must have memory units installed either in all of its slots or in none of its slots. See also memory unit.

logical segment. A single data area consisting of one or more consecutive 128-kilobyte unitary segments that is dynamically allocated by a process. The two types of logical segments are selectable segments and flat segments. See also selectable segment and flat segment.

logic board (LB). (1) See cluster switch logic board. (2) A printed wiring assembly (PWA) on which computer circuits (chips and wiring) are mounted. One type of logic board is a processor and memory board (PMB). Another type is a multifunction I/O board (MFIOB).

logical network partitioning. An HP NonStop™ TCP/IPv6 feature that allows you to divide the system into separately addressed IP subnetworks whereby applications have access only to a defined set of network interfaces (IP addresses).

logical unit number (LUN). The logical unit in the Enterprise Storage System (ESS) that maps the LDEV onto a port. Specification of the port and LUN specifies an LDEV.

login. The activity by which a user establishes a locally authenticated identity on a server. Each login has one login name.

login name. A user name associated with a session.

logon sequence. The process through which the HP NonStop™ S-series server to be managed is determined, the security constraints to interact with that server are met, and a connection with that server is established.

low-level link. A connection between the OSM or TSM client software running on a system console and the master service processors (MSPs) on an HP NonStop™ S-series server. When the HP NonStop Kernel operating system is not running, communication must take place over a low-level link. You can also communicate with a NonStop S-series server over a low-level link when the operating system is running. See also service connection.

low PIN. A process identification number (PIN) that is in the range 0 through 254. Contrast with high PIN.

LUN. See logical unit number (LUN).

MAC address. See media access control (MAC) address.

main bonding jumper. The connection between the grounded circuit conductor and the equipment grounding conductor at the service.

main memory. Data storage, specifically the chips that store the programs and data currently in use by a processor. On HP NonStop™ S-series servers, main memory is stored on the processor and memory board (PMB) in the processor multifunction

(PMF) customer-replaceable unit (CRU) and is cleared when the system is powered off.

main service entrance. The enclosure containing connection panels and switchgear, located at the point where the utility power lines enter the building.

Maintenance Interface. See Archive and Database Maintenance Interface.

maintenance switch. The Ethernet switch that links the maintenance entities in various modular components to the HP NonStop™ Open System Management (OSM) Interface.

management process. A process through which an application issues commands to a subsystem. A management process can be part of a subsystem, or it can be associated with more than one subsystem. In the latter case, the management process is logically part of each subsystem. Subsystem Control Point (SCP) is the management process for all subsystems controlled by Subsystem Control Facility (SCF).

manager. (1) For an HP NonStop™ system, the person responsible for day-to-day monitoring and maintenance tasks associated with a software subsystem on a NonStop node. (2) For a UNIX system, any person in Management and Information Services management for the site.

man page. A term sometimes used in UNIX documentation for the online or hard-copy version of a file that provides reference information. See reference page.

master service processor (MSP). A service processor (SP) that provides the basic service processor functions as well as centralized system functions such as a console port, a modem port for remote support functions, and system-load control. The enclosure containing processors 0 and 1 (group 01) also contains a pair of MSPs. See also expansion service processor (ESP).

MB. See megabyte (MB).

Measure. A tool used for monitoring the performance of the HP NonStop™ servers. Measure can be used to check the performance of a ServerNet cluster.

media access control (MAC) address. A value in the Medium Access Control sublayer of the IEEE/ISO/ANSI local area network (LAN) architecture that uniquely identifies an individual station implementing a single point of physical attachment to a LAN.

Media Interface Connector (MIC). A type of head on a fiber-optic cable that has locking wings on the sides.

megabyte (MB). A unit of measurement equal to 1,048,576 bytes (1024 kilobytes). See also gigabyte (GB), kilobyte (KB), and terabyte (TB).

member. In a structured programming language, an addressable entry within a data structure. A member can be a simple field or a data structure.

memory-exact point. A potential breakpoint location within an accelerated object file at which the values in memory (but not necessarily the values in registers) are the same as they would be if the object file were running in TNS interpreted mode or on a TNS system. Most source statement boundaries are memory-exact points. Complex statements might contain several such points: at each function call, privileged instruction, and embedded assignment. Contrast with register-exact point and nonexact point.

memory manager. An HP NonStop™ Kernel operating system process that implements the paging scheme for virtual memory. This process services requests generated by different interrupt handlers as well as by other system processes.

memory page. A unit of virtual storage. In TNS systems, a memory page contains 2048 bytes. In TNS/R systems, the page size is determined by the memory manager and can vary, depending on the processor type.

memory slot. One of eight slots for memory units on the processor and memory board (PMB). The slots are labeled MS1 through MS8.

memory unit. A unit consisting of a dual inline memory module (DIMM) or a single inline memory module (SIMM) that is installed in groups of four on the processor and memory board (PMB) of the processor multifunction (PMF) customer-replaceable unit (CRU). Memory units constitute the processor memory. The memory units in certain models of PMF CRU are not replaceable by customers or in the field.

message monitor process (MSGMON). A helper process for the ServerNet cluster monitor process (SNETMON) that runs in each processor on every node of a ServerNet cluster. MSGMON is started by the persistence manager process, $ZPM. It performs duties for SNETMON in those instances where SNETMON needs an agent in each system processor. In addition, MSGMON monitors the connections within the processor and reports changes back to SNETMON when required.

MFIOB. See multifunction I/O board (MFIOB).

MIC. See Media Interface Connector (MIC).

microcode. Any machine code or data that can run in a microprocessor. HP produces two types of microcode for HP NonStop™ S-series systems: volatile and nonvolatile. Volatile microcode is loaded into the volatile random-access memory (RAM) of some types of printed wiring assemblies (PWAs) and is not retained in a host PWA when power to the PWA is interrupted. For nonvolatile microcode, see firmware. See also millicode.

millicode. RISC instructions that implement various TNS low-level functions such as exception handling, real-time translation routines, and library routines that implement the TNS instruction set. Millicode is functionally equivalent to TNS microcode.

mirrored disk or volume. A pair of identical disk drives that are used together as a single logical volume. One drive is considered primary, and the other is called the mirror.

Each byte of data written to the primary drive is also written to the mirror drive. If the primary drive fails, the mirror drive can continue operations. See also volume.

MMF PIC. See multimode fiber-optic (MMF) plug-in card (PIC).

MMF ServerNet cable. See multimode fiber-optic (MMF) ServerNet cable.

mode. The set of attributes that specify the type and access permissions for a file. See also file mode.

modular cabinet. A standard 19-inch rack equipped with doors, bezels, and, if necessary, side panels and a power distribution unit (PDU).

modular ServerNet expansion board (MSEB). A ServerNet expansion board (SEB) that uses plug-in cards (PICs) to provide a choice of connection media for routing ServerNet packets.

module. (1) A set of components sharing a common interconnection, such as a backplane. A module is a subset of a group, and it is usually contained in an enclosure. (2) A set of I/O devices or services that share a common protocol and can be controlled by a single module driver in the extensible I/O (XIO) subsystem.

MON object type. The Subsystem Control Facility (SCF) object type for the storage pool master process.

mount. To make a fileset accessible to the users of a node.

mount point. In the Open System Services (OSS) file system, a directory that contains a mounted fileset. The mounted fileset can be in a different file system.

MRouter. A field-programmable gate array (FPGA) or application-specific integrated circuit (ASIC) that is part of the serial maintenance bus (SMB) architecture. The MRouter distributes the SMB throughout a group.

MSEB. See modular ServerNet expansion board (MSEB).

MSEB CBB. See MSEB common base board (CBB).

MSEB common base board (CBB). In modular ServerNet expansion boards (MSEBs), the printed wiring assembly (PWA) that plug-in cards (PICs) are installed on.

MSEB port. A connector on modular ServerNet expansion boards (MSEBs) used for ServerNet links. An MSEB has four fixed serial-copper ports and six plug-in card (PIC) slots that accept a variety of connection media. See also SEB port.

MSGMON. See message monitor process (MSGMON).

MSP. See master service processor (MSP).

multifunction I/O board (MFIOB). A ServerNet adapter that contains ServerNet addressable controllers (SACs) for SCSI and Ethernet; a service processor; ServerNet links to the processor, to the two ServerNet adapter slots, and to one of the ServerNet expansion board (SEB) slots; and connections to the serial maintenance bus (SMB), which connects components within an enclosure to the service processor.

multilane link. A communication link between HP NonStop™ Cluster Switches that can consist of multiple ServerNet cables. Two-lane links and four-lane links are examples of multilane links.

multimode fiber-optic (MMF) plug-in card (PIC). A plug-in card (PIC) for the modular ServerNet expansion board (MSEB) and I/O multifunction (IOMF) 2 customer-replaceable unit (CRU) that supports the multimode fiber-optic (MMF) interface.

multimode fiber-optic (MMF) ServerNet cable. A fiber-optic cable that either allows more than one mode to propagate or supports propagation of more than one mode of a given wavelength. MMF ServerNet cable typically supports shorter transmission distances than single-mode fiber-optic (SMF) ServerNet cable.

multiplexed. The action of separating data traffic from one line onto several distinct lines or of combining data traffic from several distinct lines onto one line.

mutex. See mutual exclusion (mutex).

mutual exclusion (mutex). An operating mode with interrupts disabled.

NAM. See Network Access Method (NAM).

named resource list (NRL). A table that contains one entry for each device and named process in the system. A logical device number (ldev) is an NRL index.

national-standards-body conforming POSIX.1 application. An application that both:

- Uses only the facilities described in ISO/IEC IS 9945-1:1990 and approved standards of a specific member of the ISO or IEC (the national standards body)
- Documents use of only those facilities and approved standards and documents all options and dependencies on limits

native. An adjective that can modify object code, object file, process, procedure, and mode of process execution. Native object files contain native object code, which directly uses the MIPS processor's instruction set and the corresponding conventions for register handling and procedure calls. Native processes are those created by executing native object files. Native procedures are units of native object code. Native mode execution is the state of the process when it is executing native procedures.

native link editor. See `nld` utility.

native mode. See TNS/R native mode.

native system library. Synonym for implicit library.

$NCP. The process name of the network control process.

NEC. National Electrical Code.

network. Two or more computer systems (nodes) connected so that they can exchange information and share resources. See also Expand network, wide area network (WAN), and local area network (LAN).

Network Access Method (NAM). The interface through which an Expand-over-ServerNet line-handler process communicates with the ServerNet cluster monitor process (SNETMON).

network control process. A process pair, named $NCP, that runs in each system of an Expand network. $NCP establishes and terminates system-to-system connections, maintains network-related system tables (including the network routing table, NRT), calculates the most efficient way to transmit data to other systems in the network, monitors and logs changes in the status of the network and its systems, informs $NCPs at neighbor systems of changes in line or Expand line-handler process status, and aborts pending requests when all paths go down. See also network routing table (NRT).

Network Information Service (NIS). A distributed name service (formerly known as Yellow Pages) developed by Sun Microsystems. See also Domain Name System (DNS).

network routing table (NRT). A table that resides in each processor in each system in a network. The NRT associates each destination system with the logical device (LDEV) number of the best-path route Expand line-handler process to use to send messages to that system. See also network control process.

network topology. The physical layout of components that define a system, a local area network (LAN), or a wide area network (WAN).

neutral. (1) The conductor used as the primary return for current during normal operation of electrical equipment. (2) The junction of the legs in a wye circuit. See also wye.

NIS. See Network Information Service (NIS).

`nld` utility. (1) The utility that collects, links, and modifies code and data blocks from one or more object files to produce a target TNS/R native object file. The `nld` utility is similar to the Binder program used in the TNS development environment. (2) The native link editor invoked during system generation to build the TSYSCLR and TSYSDP2 files.

NNA. See node-numbering agent (NNA).

NNA PIC. See node-numbering agent (NNA) plug-in card (PIC).

node. (1) A uniquely identified computer system connected to one or more other computer systems in a network. See also Expand node and ServerNet node. (2) An endpoint in a ServerNet fabric, such as a processor or ServerNet addressable controller (SAC).

node number. A number used to identify a member system in a network. The node number is usually unique for each system in the network. See also node and ServerNet node number.

node-numbering agent (NNA). A field-programmable gate array (FPGA) in a single-mode fiber-optic (SMF) plug-in card (PIC) that translates the node number of each ServerNet packet entering or exiting the external ServerNet fabrics.

node-numbering agent (NNA) plug-in card (PIC). A plug-in card (PIC) for the modular ServerNet expansion board (MSEB) that supports the node-numbering agent (NNA) interface.

node routing ID. See ServerNet node routing ID.

noncanonical input mode. For an Open System Services (OSS) process, a terminal input mode in which data is made available to the process when a timer expires or when a certain number of characters have been entered. Noncanonical data is not grouped into logical lines of input. This mode is sometimes called block mode or transparent mode. Contrast with canonical input mode.

nonclustered. Lacking the quality of belonging to a cluster.

nonconfigured object. An object that comes into existence after Subsystem Control Facility (SCF) is running and that was created in response to activity outside the SCF environment. An SCF STATUS command can display the name of a nonconfigured object, but its state is UNKNOWN.

nondedicated (public) LAN. A local area network (LAN) connected to the Ethernet ports on an Ethernet 4 ServerNet adapter (E4SA), Fast Ethernet ServerNet adapter (FESA), or Gigabit Ethernet ServerNet adapter (GESA). Unlike a dedicated service LAN, a public LAN supports the connection of many types of servers and workstations. System consoles can be connected to a public LAN, but such system consoles cannot use all the OSM or TSM client applications. See also dedicated service LAN.

nonessential firmware. Code that is used for support routines such as self-test diagnostics and that can be overwritten during flash programming without affecting the next power-up operation. Contrast with essential firmware.

nonexact point. A code location within an accelerated object file that is between memory-exact points. The mapping between the TNS program counter and corresponding RISC instructions is only approximate at nonexact points, and interim changes to memory might have been completed out of order. Breakpoints cannot be applied at nonexact points. Contrast with memory-exact point and register-exact point.

noninteractive mode. A mode of operation that usually involves a command file (an EDIT file that contains a series of commands). Contrast with interactive mode.

nonlinear load. Electrical load for which the instantaneous current is not proportional to the instantaneous voltage. Consequently, the local impedance varies with the voltage.

nonsensitive command. A command that can be issued by any user or program that is allowed access to a subsystem—that is, a command on which the subsystem imposes no further security restrictions. For Subsystem Control Facility (SCF), nonsensitive commands are those that cannot change the state or configuration of objects; most of them are information commands. Contrast with sensitive command.

NonStop™ Cluster Switch. See HP NonStop™ Cluster Switch (model 6770).

NonStop™ zone. A branch of the power-distribution system that provides power directly to HP NonStop computer equipment.

NonStop™ Kernel operating system. See HP NonStop™ Kernel operating system.

NonStop™ ServerNet Cluster. See HP NonStop™ ServerNet Cluster (ServerNet Cluster).

NonStop™ ServerNet Switch. See HP NonStop™ ServerNet Switch (model 6780).

NonStop™ TCP/IP. See HP NonStop™ TCP/IP.

NonStop™ TCP/IP process. See HP NonStop™ TCP/IP process.

NonStop™ TCP/IP subsystem. See HP NonStop™ TCP/IP subsystem.

NonStop™ TCP/IPv6. See HP NonStop™ TCP/IPv6.

NonStop Technical Library. See HP NonStop™ Technical Library (NTL).

normal mode. Electromagnetic interference that occurs between current-carrying conductors (for example, line to neutral).

normal mode

notification. Another name for an incident report created by the OSM and TSM Notification Director. When incident reports are dialed out to service providers, this process is also referred to as remote notification.

NRL. See named resource list (NRL).

NRT. See network routing table (NRT).

NSR-D processor. See HP NonStop™ System RISC Model D processor (NSR-D processor).

NSR-E processor. See HP NonStop™ System RISC Model E processor (NSR-E processor).

NSR-G processor. See HP NonStop™ System RISC Model G processor (NSR-G processor).

NSR-T processor. See HP NonStop™ System RISC Model T processor (NSR-T processor).

NSR-V processor. See HP NonStop™ System RISC Model V processor (NSR-V processor).

NSR-W processor. See HP NonStop™ System RISC Model W processor (NSR-W processor).

NSR-X processor. See HP NonStop™ System RISC Model X processor (NSR-X processor).

NSR-Y processor. See HP NonStop™ System RISC Model Y processor (NSR-Y processor).

null object type. A placeholder object type for the Subsystem Control Facility (SCF) NAMES and VERSION commands, which do not require explicit specification of a particular object type.

null string. In C and C++ programs, a character string that begins with a null character. This term is synonymous with *empty string*.

OBEY file. See command file.

object. One or more of the devices, lines, processes, and files in a subsystem; any entity subject to independent reference or control by one or more subsystems. In the Subsystem Control Facility (SCF), each object has an object type and an object name.

object-code library. Synonym for library.

object code file. A file containing compiled machine instructions for one or more routines. This file can be an executable loadfile for a program or library or a not-yet-executable linkfile for some program module. On other systems, an object code file is also known as a *binary* or as an *executable*.

object name. A unique name for a Subsystem Control Facility (SCF) object within a subsystem.

object-name template. A name that stands for more than one Subsystem Control Facility (SCF) object. Such a name includes one or more wild-card characters, such as * (asterisk) and ? (question mark). See also wild-card character.

object type. The category of Subsystem Control Facility (SCF) objects to which a specific SCF object belongs. For example, a specific disk has the object type DISK, and a specific terminal can have the object type SU. Each subsystem has a set of object types for the objects it manages.

obsolescent. An indication that a feature or facility exists for compatibility with older versions or drafts of a standard. Obsolescent features or facilities should not be used because they might be removed from future versions of a standard and therefore might not be portable.

offline. (1) Used to describe tasks that are performed outside of the control of an application or computer system. (2) Used to describe tasks that require system resources to be shut down. Contrast with online.

offline change. Any change that requires system resources to be shut down. Offline changes are usually performed during a planned outage. Contrast with online change.

offline configuration. Configuration performed offline by SYSGENR. If necessary, you edit the CONFTEXT configuration file to create a new configuration and then run the Distributed Systems Management/Software Configuration Manager (DSM/SCM) (which in turn runs SYSGENR) to generate a system image for the new configuration.

ohm. The standard unit for measuring resistance.

online. Used to describe tasks that can be performed while the HP NonStop™ Kernel operating system and system utilities are operational. Contrast with offline.

online change. Any change that can be performed while an application or its system resources are operational. In some situations, online changes might temporarily affect subsystem and application availability. For example, altering the characteristics of a communications line might temporarily affect applications that use the communications line. Contrast with offline change.

Online Support Center (OSC). The group of support specialists within the HP Global Customer Support Center (GCSC) who respond to telephone calls regarding system problems and diagnose malfunctioning systems using remote diagnostic links.

open file. In the Open System Services (OSS) file system, a file with a file descriptor.

open file description. In the Open System Services (OSS) file system, a data structure within an HP NonStop™ node that contains information about the access of a process or of a group of processes to a file. An open file description records such attributes as the file offset, file status, and file access modes. An open file description is associated with only one open file but can be associated with one or more file descriptors.

open migration. In the Open System Services (OSS) file system, the set of events and outcomes that occur when an open file description is inherited by a child process in a different processor than its parent process. Contrast with open propagation.

open propagation. In the Open System Services (OSS) file system, the set of events and outcomes that occur when an open file description is inherited by a child process in the same processor as its parent process. Contrast with open migration.

Open SCSI. A subsystem that provides the hardware and software for a SCSI-2 open interface that runs on HP NonStop™ S-series servers and to which developers can attach small computer system interface (SCSI) devices.

open system. A system with interfaces that conform to international computing standards and therefore appear the same regardless of the system's manufacturer. For example, the Open System Services (OSS) environment on HP NonStop™ systems conforms to international standards such as ISO/IEC IS 9945-1:1990 (ANSI/IEEE Std. 1003.1-1990, also known as POSIX.1), national standards such as FIPS 151-2, and portions of industry specifications such as the X/Open Portability Guide Version 4 (XPG4).

Open System Services (OSS). An open system environment available for interactive or programmatic use with the HP NonStop™ Kernel operating system. Processes that run in the OSS environment usually use the OSS application program interface. Interactive users of the OSS environment usually use the OSS shell for their command interpreter. Synonymous with *Open System Services (OSS) environment*. Contrast with Guardian.

Open System Services (OSS) environment. The HP NonStop™ Kernel Open System Services (OSS) application program interface (API), tools, and utilities.

Open System Services (OSS) Monitor. A Guardian utility that accepts commands affecting OSS objects through an interactive Guardian interface named the Subsystem Control Facility (SCF).

Open System Services (OSS) signal. A signal model defined in the POSIX.1 specification and available to TNS processes and TNS/R native processes in the OSS environment. OSS signals can be sent between processes.

Open Systems Interconnection (OSI). A seven-layer network architecture model defined by the International Organization for Standardization (ISO). The two lowest layers deal with the physical connections and their protocols. The five upper layers deal with network services, such as network file transfers and access to remote databases.

Open Systems Interconnection Layer 2. The data-link control level of the Open Systems Interconnection (OSI) model, composed of asynchronous or minimal line control protocols, byte-oriented or character-oriented protocols, and bit-synchronous or bit-oriented protocols. Data link protocols can be defined in terms of method of access of data, link relationship of stations, error detection scheme, error recovery, message formatting, logical half-duplex or full-duplex operation, code, and machine transparency.

operating system image. See OSIMAGE.

operational environment. The conditions under which your system performs. These conditions include the devices and communications lines that are made active and the system and application processes that are started at system startup.

operator. (1) A symbol, such as an arithmetic or conditional operator, that performs a specific operation on operands. (2) In Network Control Language (NCL), a lexical element used for working on terms in expressions. The five types of operators are parenthetical, arithmetic, Boolean, relational, and string. (3) For an HP NonStop™ system, the person or program responsible for day-to-day monitoring and maintenance tasks associated with the HP NonStop Kernel operating system and the hardware of a NonStop node. The operator issues commands to subsystems; retrieves, examines, and responds to event messages; or does any combination of those things. See also local operator. Contrast with administrator. (4) For a UNIX system, any interactive user of that system.

operator message. A message, intended for an operator, that describes a significant event on an HP NonStop™ S-series system. An operator message is the displayed-text form of an Event Management Service (EMS) event message.

option. In a UNIX or Open System Services (OSS) command, a flag and its parameters or a flag without parameters.

ordinary library. A dynamic-link library (DLL) or shared run-time library (SRL) that is not public. See ordinary dynamic-link library (ordinary DLL).

ordinary dynamic-link library (ordinary DLL). A dynamic-link library (DLL) or shared run-time library (SRL) that is not public. The code file is found at run time and can be provided by the user. Contrast with public dynamic-link library (public DLL).

orphaned process group. In the Open System Services (OSS) environment, a process group in which the parent of every member is also either a member of the process group or a member of a different session.

orphan file. In the Open System Services (OSS) environment, a file with no corresponding inode in the PXINODE file.

orphan inode. In the Open System Services (OSS) environment, an inode that appears in the PXINODE file but has no links in the PXLINK file.

OSC. See Online Support Center (OSC).

OSCONFIG file. In G-series release version updates (RVUs), a configuration file built during system generation that contains only Software Problem Isolation and Fix Facility (SPIFF) and Software Identification (SWID) tool records. In D-series and earlier RVUs, the Configuration Utility Program (COUP) uses the $SYSTEM.SYS*nn*.OSCONFIG file to store its configuration information.

OSI. See Open Systems Interconnection (OSI).

OSIMAGE. A file built during system generation that contains the complete image of the HP NonStop™ Kernel operating system that runs in each processor in the system.

OSM. See HP NonStop™ Open System Management (OSM) Interface.

OSM Event Viewer. OSM replacement for the TSM Event Viewer.

OSM Low-Level Link. OSM replacement for the TSM Low-Level Link.

OSM Notification Director. OSM replacement for the TSM Notification Director.

OSM Service Connection. OSM replacement for the TSM Service Application.

OSS. See Open System Services (OSS).

OSS environment. See Open System Services (OSS) environment.

OSS Monitor. See Open System Services (OSS) Monitor.

OSS process ID. In the Open System Services (OSS) environment, the unique identifier that identifies a process during the lifetime of the process and during the lifetime of the process group of that process. See also PID.

OSS signal. See Open System Services (OSS) signal.

OSS user ID. See HP NonStop™ Kernel user ID.

outage. Time during which a computer system is not capable of doing useful work. Outages can be planned or unplanned. From the end user's perspective, an outage is any time an application being used is not available. See also planned outage and unplanned outage.

outage minutes. A metric for measuring outages that translates percentages into minutes/year of downtime.

output destination. The resource to which Subsystem Control Facility (SCF) sends its responses to commands. SCF can direct output to a disk file, an application process, a terminal, or a printer. The initial output destination is determined by the form of the RUN command used to initiate SCF. The output destination can be changed dynamically during an SCF session.

owner. (1) In the case of a disk file, the user or program that created the file, or a user or program to whom the creator has given the file with the File Utility Program (FUP) GIVE command. (2) In the case of a process, the user or program that created the process or, if the PROGID option was specified in the FUP SECURE command for the code file, the user or program that owns the code file. (3) In the case of a token or other definition, the subsystem that provided the definition. (4) In the case of a subsystem, the company or organization that provides the subsystem, or the 8-character string identifying that company.

packet. A block of information that contains fields for addressing, sequencing information, possible priority indicators, and a portion of a message or an entire message. See also ServerNet packet.

page. See memory page.

Parallel Library TCP/IP. An HP product that provides increased performance and scalability over conventional Transmission Control Protocol/Internet Protocol (TCP/IP). Parallel Library TCP/IP coexists with conventional TCP/IP on HP NonStop™ S-series systems and supports Ethernet 4 ServerNet adapters (E4SAs), Fast Ethernet ServerNet adapters (FESAs), Gigabit Ethernet ServerNet adapters (GESAs), and ServerNet wide area network (SWAN) concentrators. See also HP NonStop™ TCP/IP.

PARAM. An HP Tandem Advanced Command Language (TACL) command and a Subsystem Control Facility (SCF) command you can use to create a parameter and give it a value. The TACL process stores the values of parameters assigned by the PARAM command and sends the values to applications that request parameter values.

parent directory. A particular directory in the hierarchy of directories within a file system. The parent directory for a directory contains an entry for that specific directory and is identified in that directory as the directory immediately above it in the hierarchy. The parent directory for a file contains an entry for that file.

parent process. The process that created a given process or (if the creating process has stopped) a process that has inherited a given process. See also child process.

parent process ID. In the Open System Services (OSS) environment, an attribute of a child process determined by the parent process. The parent process ID is the OSS process ID of the current parent process.

passthrough terminator. See SCSI passthrough terminator.

path. The route between a processor and a subsystem. If a subsystem is configured for fault tolerance, it has a primary path (from the primary processor) and a backup path (from the backup processor).

pathname. In the Open System Services (OSS) file system and Network File System (NFS), the string of characters that uniquely identifies a file within its file system. A pathname can be either relative or absolute. See also ISO/IEC IS 9945-1:1990 (ANSI/IEEE Std. 1003.1-1990 or POSIX.1), Clause 2.2.2.57.

pathname component. See filename.

pathname resolution. In the Open System Services (OSS) environment, the process of associating a single file with a specified pathname.

pathname-variable limits. Limits that can vary within the Open System Services (OSS) file hierarchy; that is, the limits on a pathname variable that can vary according to the directory in which pathname resolution begins.

path prefix. In the Open System Services (OSS) environment, a pathname with an optional final slash (/) character that refers to a directory.

PDC. See phase-loss detector/contactor (PDC).

PDP. See power distribution panel (PDP).

PDU. See computer-room power center (CRPC) and power distribution unit (PDU).

peak load current. The maximum instantaneous load over a designated interval of time.

PEEK. A utility program that reports statistics on resource use in a processor. PEEK is used to ensure proper allocation of memory and processes in a system after system load.

peer fabric. The fabric on which an operation is not taking place. The X and Y fabrics are peers. If an action is being performed on one fabric, the other fabric is the peer fabric.

peer service processors. A pair of service processors (X and Y) in a service processor (SP) domain. Peer service processors function similarly to a fault-tolerant process pair in an HP NonStop™ K-series system. See also service processor (SP).

pending incident report. An incident report that has never been delivered to your service provider, either because delivery to both the primary and backup dial-out points was unsuccessful or because the incident report was generated at an unattended site.

pending signal. A signal that has been generated for a process but has not been delivered. Pending signals are usually blocked signals.

periodic incident report. A type of incident report that is generated periodically to test the connection to the service provider and report the current system configuration. The default frequency is 20 days.

peripheral enclosure. An enclosure that contains components related to one or more peripherals. The 519x tape subsystem is an example of a peripheral enclosure. Peripheral enclosures are not part of the set of system enclosures. Contrast with system enclosure.

Peripheral Utility Program (PUP). A utility used in D-series and earlier release version updates (RVUs) to manage disks and other peripheral devices. In G-series RVUs, similar functions are performed by the Subsystem Control Facility (SCF).

persistence. For the Subsystem Control Facility (SCF), the capability of a generic process to restart automatically if it was stopped abnormally. You configure this capability by specifying a nonzero AUTORESTART value in an ADD command.

persistence count. The number of times the $ZPM persistence manager process will restart a generic process that has been terminated abnormally. A generic process with an AUTORESTART value of 10 (the maximum) is said to have a persistence count of 10. See also persistence.

persistence manager process. The $ZPM process that is started and managed by the $ZCNF configuration utility process and that starts generic processes in G-series release version updates (RVUs) and manages their persistence.

persistent configuration. A configuration that remains the same from one system load to another.

persistent process. A process that must always be either waiting, ready, or executing. Persistent processes are usually controlled by a monitor process that checks on the status of persistent processes and restarts them if necessary.

phase-loss detector/contactor (PDC). Equipment used to detect the interruption (for 50 milliseconds or longer) or the complete loss of one or more phases of power to computer equipment. Upon detection of a phase dropout, the contactor shuts down all input phases to the system equipment, thereby allowing smooth system shutdown and recovery.

physical interface (PIF). The hardware components that connect a system node to a network.

physical link interfaces. Communications standards defined by standards organizations. The following physical link interfaces are supported for the ServerNet wide area network (SWAN) concentrator: RS-232, RS-442, RS-449, V.35, and X.21.

Physical view. One of several views of a server available in the view pane of the Management window of the OSM Service Connection, TSM Service Application, and OSM and TSM Low-Level Link. A Physical view of a server is a view of all the enclosures and is intended to represent the actual floor plan at the site. A Physical view of an enclosure is a visual representation of the physical placement of supported resources inside the enclosure. See also Connection view.

PIB. See power interface board (PIB).

PIC. See plug-in card (PIC).

PID. In the Open System Services (OSS) environment, a synonym for process ID. OSS process ID is the preferred term in HP NonStop™ S-series system publications.

In the Guardian environment, PID is sometimes used to mean either:

- A Guardian process identifier such as the process ID
- The *cpu, pin* value that is unique to a process within a node (See HP NonStop™ Kernel user ID.)

PIF. See physical interface (PIF).

PIN. See process identification number (PIN).

ping. A utility used to verify connections to one or more remote hosts. The ping utility uses the Internet control message protocol (ICMP) echo request and echo reply packets to determine whether a particular IP system on a network is functional. The ping utility is useful for diagnosing IP network or router failures.

pipe. In the Open System Services (OSS) environment, an unnamed FIFO, created programmatically by invoking the `pipe()` function or interactively with the shell pipe syntax character (`|`). A shell pipe redirects the standard output of one process to become the standard input of another process. A programmatic pipe is an interprocess communication mechanism.

planned outage. Time during which a computer system is not capable of doing useful work because of a planned interruption. A planned outage can be time when the system or user application is shut down to allow for servicing, upgrades, backup, or general maintenance.

planner. The Distributed Systems Management/Software Configuration Manager (DSM/SCM) user who is responsible for planning and managing new software revisions. The planner uses the DSM/SCM Planner Interface to carry out these functions.

Planner Interface. A graphical user interface (GUI) to the Distributed Systems Management/Software Configuration Manager (DSM/SCM) that runs on the host system. It provides an interface to all the host DSM/SCM planner functions.

plug-in card (PIC). A replaceable component that provides a unique function when installed in a customer-replaceable unit (CRU) or field-replaceable unit (FRU). PICs for modular ServerNet expansion boards (MSEBs) and I/O multifunction (IOMF) 2 CRUs provide a choice of connection media for attaching ServerNet cables.

PM. See product module (PM).

PMB. See processor and memory board (PMB).

PMCU. See power monitor and control unit (PMCU).

PMF CRU. See processor multifunction (PMF) CRU.

PMF 2 CRU. See processor multifunction (PMF) 2 CRU.

pNA+. The support for the Transmission Control Protocol/Internet Protocol (TCP/IP) layers and the Ethernet interface provided by Integrated Systems Inc. as part of the Portable Silicon Operating System (pSOS) system product. pNA+ is provided as part of the wide area network (WAN) architecture in each ServerNet wide area network (SWAN) concentrator communications line interface processor (CLIP).

Point-to-Point Protocol (PPP). A data communications protocol that provides a standard method of encapsulating Transmission Control Protocol/Internet Protocol (TCP/IP) information over point-to-point links. OSM and TSM use PPP to provide TCP/IP communication over a dial-up connection.

POOL object type. The Subsystem Control Facility (SCF) object type for storage pool storage pools.

port. (1) A data channel that connects to other devices or computers. (2) A connector to which a cable can be attached. The system transmits and receives data or requests through ports on ServerNet adapters and processor multifunction (PMF) customer-replaceable units (CRUs). A port is also called a connector. (3) The entrance or physical access point (such as a connector) to a computer, multiplexer, device, or network where signals are supplied, extracted, or observed.

portable application. An application that can execute on a wide range of hardware systems from multiple manufacturers. A portable application is a program that can be moved with little or no change in its source code from another manufacturer's system to an HP NonStop™ system.

portable filename character set. The set of characters that includes the Roman uppercase and lowercase letters, the Arabic numerals, the period, the underscore, and the hyphen. The hyphen cannot be the first character of a portable filename.

portable pathname character set. The set of characters that includes the Roman uppercase and lowercase letters, the Arabic numerals, the period, the underscore, the slash (/), and the hyphen. The hyphen cannot be the first character of a portable pathname.

Portable Silicon Operating System (pSOS) system product. A product of Integrated Systems Inc. that provides support for industry-standard communications protocols based on the UNIX operating system. It is used as a compact multitasking kernel operating system for PowerPCs and similar systems.

position ID. A character that indicates the position an HP NonStop™ Cluster Switch occupies in a network topology. The position ID is a component of the two-character cluster switch name. The cluster switch name includes an external fabric ID (X or Y) as the first character and a position ID as the second character. For example, the cluster switch name X3 indicates that the cluster switch serves the external ServerNet X fabric and occupies position 3 in the topology. Supported values for position IDs are 1 through 9 or A through Z. Currently supported topologies (star, split-star, and tri-star) use position IDs 1, 2, and 3.

position-independent code (PIC). Executable program or library code that is designed to be loaded and executed at any virtual memory address, without any modification. Addresses that can be modified by the loader do not appear in PIC code, only in data that can be modified by the loader. See also dynamic-link library (DLL).

POSIX. The Portable Operating System Interface, as defined by the Institute of Electrical and Electronics Engineers (IEEE) and the American National Standards Institute (ANSI). Each POSIX interface is separately defined in a numbered ANSI/IEEE standard or draft standard. The application program interface (API), known as POSIX.1, has become ISO/IEC IS 9945-1:1990.

power distribution panel (PDP). A group of panel assemblies that composes a single panel that includes buses and overcurrent protection devices (with or without switches). A PDP is used for the control of power circuits.

power distribution unit (PDU). An assembly of outlets and fuses that can be installed on both sides of a modular cabinet. This assembly provides power and over-current protection for the components installed in the modular cabinet. See also computer-room power center (CRPC).

power domain. A set of customer-replaceable units (CRUs) and field-replaceable units (FRUs) that share a set of power rails. For telco Central Office (CO) systems, the power domain is the entire system.

power factor. The ratio of real power to apparent power (that is, kilowatts/kilovoltamperes). The power factor for a sinusoidal load is determined by the position of the applied voltage waveform with respect to the current drawn by the load. When voltage and current are in phase with each other, the power factor is unity, and the power for the load is equal to the product of the applied voltage and load current (P=EI). When the current waveform lags the voltage waveform, the load is inductive. Conversely, when the current waveform leads the voltage waveform, the load is capacitive. In either case, the power for the load is equal to the product of the applied voltage, load current, and the angular displacement between the voltage and current waveforms (P=EIcosf). Nonlinear (nonsinusoidal) loads also have a power factor. However, the power factor for a nonsinusoidal load reflects harmonic content and not angular displacement.

power factor correction. The addition of a reactive component to offset the angular displacement of a sinusoidal load. Traditionally, the normal power factor for a facility is inductive, so the normal correction involves the addition of capacitors to offset the lagging power factor. The capacitors offset part or all of the inductive reactance, making the total circuit more nearly in phase with the applied voltage. The power factor for nonlinear (nonsinusoidal) loads cannot be corrected through the addition of simple reactive components. Harmonic filters are required to correct the power factor of nonlinear loads.

power interface board (PIB). In system enclosures with power shelves, a board mounted on the bulkhead located behind the power supplies in the power shelf. The PIB provides electrical connection between the power supplies and DC power cables.

power monitor and control unit (PMCU). A field-replaceable unit (FRU) that connects the batteries to the DC power distribution bus in an HP NonStop™ S-series enclosure and provides a means of disconnecting the batteries for powering off the system. The PMCU also provides a means for the service processor (SP) to diagnose the condition of the batteries, fans, and power supplies; to regulate the voltage supplied to the fans; and to provide the interface to the group ID switches and service light-emitting diodes (LEDs). A group contains two PMCUs, one for each of the two DC power distribution buses.

power shelf. In HP NonStop™ S7400, S7600, and Sxx000 processor enclosures and I/O enclosures containing I/O multifunction (IOMF) 2 customer-replaceable units (CRUs), an assembly residing below the chassis consisting of power supplies and supporting circuitry that provides DC power to the enclosure.

power supply. (1) In system enclosures without power shelves, the component on the processor multifunction (PMF) customer-replaceable unit (CRU) or the I/O multifunction (IOMF) CRU that converts standard AC line voltage into the DC voltages needed by the group components in the enclosure. (2) In system enclosures with power shelves, the component located in the power shelf that converts standard AC line voltage into DC voltage and delivers it to the PMF CRUs or IOMF CRUs in that enclosure, which in turn supply the DC voltages needed by the group components in the enclosure.

PPL. See process-pair list (PPL).

PPP. See Point-to-Point Protocol (PPP).

preemption. A form of late binding in which a symbolic reference to a symbol defined in the same dynamic-link library is instead bound to a definition in another loadfile.

preferences file. A file that contains configuration information for the graphical user interface (GUI) portion of the OSM and TSM client software. The preferences file is used by the OSM and TSM client software at system startup.

preferred path. See primary path.

preprocessing commands. Commands specifying unique run-time parameters that can override your default system parameters. These commands can assign process file names, select backup media formats, and define utility options during system configuration.

preset. A linker operation that sets the correct values (addresses) of imported symbols according to the environment seen by the linker. If the loader encounters the same environment at load time, it avoids adjusting these values, which reduces loading overhead. (See fastLoad.) If not, the loader resets these values to match the load-time environment.

primary path. A path enabled as the preferred path. When a primary path is disabled, an alternate path becomes the primary path.

primary processor. The processor that is designated as *owning* the ServerNet addressable controller (SAC) connected to separate processors running the HP NonStop™ Kernel operating system. The primary processor is the processor that has direct control over the SAC. Contrast with backup processor.

private dynamic-link library (private DLL). See ordinary dynamic-link library (ordinary DLL).

problem incident report. A type of incident report that reports a problem in the server. A problem incident report is generated when changes occur on the server that could directly affect the availability of system resources.

procedure entry-point (PEP) table. A table in a TNS object file that contains the entry point addresses for each procedure and is located in the first page of each code segment.

process. (1) A program that has been submitted to the operating system for execution, or a program that is currently running in the computer. (2) An address space, a single thread of control that executes within that address space, and the system resources required by that thread of control.

process group. In the Open System Services (OSS) environment, a set of processes that can signal associated processes. Each process in a node is a member of a process group. The process group has a process group ID. A new process becomes a member of the process group of its creator.

process group ID. In the Open System Services (OSS) environment, the unique identifier representing a process group during its lifetime.

process group leader. In the Open System Services (OSS) environment, the process that has the process group ID of its process group as its OSS process ID.

process group lifetime. In the Open System Services (OSS) environment, the period that begins when a process group is created and ends when the lifetime of the last remaining process of the group ends.

process ID. In the Guardian environment, the content of a 4-integer array that uniquely identifies a process during the lifetime of the process. See also PID.

process identification number (PIN). A number that uniquely identifies a process running in a processor. The same number can exist in other processors in the same system. See also process ID.

process image file. On a UNIX system, an executable object file. In some Guardian product externals and end-user publications, an executable object file is referred to as a *program file*. See also object code file.

process lifetime. The period that begins when an Open System Services (OSS) process is created and ends when its OSS process ID is returned to the system for reuse.

PROCESS object type. In a subsystem, the object type for the subsystem manager process itself or any generic process.

process-pair list (PPL). A table that contains one entry for each process or process pair associated with a named resource. For some devices, a single PPL entry is associated with multiple named resource list (NRL) entries. For process pairs, a single PPL entry describes both members of the pair.

processor. (1) A functional unit of a computer that reads program instructions, moves data between processor memory and the input/output controllers, and performs arithmetic operations. A processor is sometimes referred to as a central processing unit (CPU), but HP NonStop™ servers have multiple cooperating processors rather than a single CPU. (2) One or more computer chips, typically mounted on a logic board, that are designed to perform data processing or to manage a particular aspect of computer operations.

processor and memory board (PMB). A logic board that has lockstepped microprocessors, the main memory system, and the ServerNet memory interface (SMI) application-specific integrated circuits (ASICs) to act as an interface between the microprocessors and memory and the ServerNet fabrics. This board is part of the processor multifunction (PMF) customer-replaceable unit (CRU).

processor cache. A small, fast memory holding recently accessed data in order to speed up subsequent access to the same data. Cache memory is built from faster memory chips than main memory. It is most often used with process or main memory but is also used in network data transfer (to maintain a local copy of data) and so forth.

processor dump. A copy of the memory of a processor. A dump can be to disk or to tape. See also ServerNet dump and tape dump.

processor enclosure. An HP NonStop™ S-series system enclosure containing one group, which includes processors, ServerNet adapters, disk drives, components related to the ServerNet fabrics, and components related to electrical power and cooling for the enclosure.

processor multifunction (PMF) CRU. (1) An HP NonStop™ S-series customer-replaceable unit (CRU) that contains a power supply, service processor (SP), ServerNet router 1, Ethernet controller, three ServerNet addressable controllers (SACs), and a processor and memory system in a single unit. The PMF CRU consists of three subassemblies: the processor and memory board (PMB), the multifunction I/O board (MFIOB), and the power supply subassembly. (2) A collective term for both PMF CRUs and PMF 2 CRUs when a distinction between the two types of CRUs is not required.

processor multifunction (PMF) 2 CRU. An HP NonStop™ S-series customer-replaceable unit (CRU) that contains a power supply, service processor (SP), ServerNet router 2, Ethernet controller, three ServerNet addressable controllers (SACs), and a processor and memory system in a single unit. The PMF 2 CRU consists of three subassemblies: the processor and memory board (PMB), the multifunction I/O board (MFIOB), and the power supply subassembly.

product module (PM). The part of the Subsystem Control Facility (SCF) subsystem that is responsible for subsystem-specific command processing.

profile. Default values used by the Distributed Systems Management/Software Configuration Manager (DSM/SCM) when processing requests. The three types of profiles are the Configuration Manager profile, the system profile, and the target profile.

PROFILE object type. The Subsystem Control Facility (SCF) object type for the storage subsystem configuration profile.

program. See program file.

program file. An executable object code file containing a program's main routine plus related routines statically linked together and combined into the same object file. Other routines shared with other programs might be located in separately loaded libraries. A program file can be named on a RUN command; other code files cannot. See also object code file.

pSOS system product. See Portable Silicon Operating System (pSOS) system product.

public dynamic-link library (public DLL). Optional native-mode executable code modules available to all native user processes. A public library that is specified in the public library registry, supplied by HP or, optionally, a user. Contrast with ordinary dynamic-link library (ordinary DLL).

public LAN. See nondedicated (public) LAN.

public library. A dynamic-link library (DLL) or shared run-time library (SRL) that is known to the operating system, available for execution by any process or user, and is not an implicit library.

public shared run-time library (public SRL). A TNS/R library supplied by HP.

PUP. See Peripheral Utility Program (PUP).

quad-integrated communications controller (QUICC). The Motorola MC68360 chip. For HP NonStop™ S-series servers, the QUICC is used as the service processor (SP) and is the main part of the ServerNet wide area network (SWAN) concentrator communications line interface processor (CLIP).

quality power. The attributes and configuration of the power-distribution systems installed within a facility that best serve the power needs of that facility's electrical equipment (for example, computer systems, air conditioning, and so on), providing the minimum possible disruption to equipment operation.

QUICC. See quad-integrated communications controller (QUICC).

R1. See ServerNet router 1.

R2. See ServerNet router 2.

raceway. An enclosed channel used to hold wires, cables, or busbars. Most raceways have removable tops to facilitate the installation or removal of their contents.

rack. A structure that houses a chassis, power shelf, and other system components. The HP NonStop™ S-series server is designed to be mounted in an industry-standard 19-inch rack or a NonStop S-series frame. See also frame.

radio frequency interference (RFI). Forms of conducted or radiated interference that might appear in a facility as either normal or common-mode signals. The frequency of the interference can range from the kilohertz to gigahertz range. However, the most troublesome interference signals are usually found in the kilohertz to low megahertz range. At present, the terms radio frequency interference and electromagnetic interference (EMI) are usually used interchangeably.

range of servers. See HP NonStop™ servers.

read-only file system. A file system with implementation-defined characteristics that restrict changes to the files within that file system.

read/write head. An electromagnet that can pick up (read) electronic pulses and record (write) electronic pulses on a magnetic disk or tape. The electronic pulses are interpreted by the processor as binary data. See also disk drive and tape drive.

real group ID. An attribute of an Open System Services (OSS) process. When an OSS process is created, the real group ID identifies the group of the user or parent process that created the process. The real group ID can be changed after process creation.

real user ID. An attribute of an Open System Services (OSS) process. When an OSS process is created, the real user ID identifies the user or parent process that created the process. The real user ID can be changed after process creation.

$RECEIVE. The name of a file through which a process receives and optionally replies to messages from other processes.

reconfiguration. The act of changing the hardware or software configuration of a running system. Examples include installing a new software release version update (RVU), adding hardware peripherals, and restructuring a database. Reconfiguring a system might or might not require a planned outage.

reduced instruction-set computing (RISC). A processor architecture based on a relatively small and simple instruction set, a large number of general-purpose registers, and an optimized instruction pipeline that supports high-performance instruction execution. Contrast with complex instruction-set computing (CISC).

reexported library. A library whose symbols are made available by another dynamic-link library (DLL) to any localized client of that DLL. Reexport is an attribute of the DLL's libList entry for that library. This attribute is specified by the DLL's programmer and recorded by the linker as a DLL is built. It affects only localized clients of the DLL. This feature allows a symbol to be moved from one DLL to another without relinking clients of the original DLL.

Reexporting is transitive; that is, if A reexports B and B reexports C, then A reexports C. Reexported libraries can reexport other libraries to form a succession of reexported libraries of arbitrary length.

reference page. In Open System Services (OSS) and Distributed Computing Environment (DCE), the online or hard-copy version of a file that provides reference information for a software facility. Some UNIX product externals and end-user publications use the term *man page* instead, referring either to the online delivery mechanism used to display the file (usually the shell `man` command) or to the nature of the file as part of a publication.

register-exact point. A synchronization location within an accelerated object file at which both of these statements are true:

- All live TNS registers plus all values in memory are the same as they would be if the object file were running in TNS mode or TNS interpreted mode or on a TNS system.
- All accelerator code optimizations are ended.

Register-exact points are a small subset of all memory-exact points. Procedure entry and exit locations and call-return sites are usually register-exact points. All places where the program might switch into or from TNS mode or TNS interpreted mode are register-exact points. Contrast with memory-exact point and nonexact point.

regular file. In the Open System Services (OSS) file system, a file that is a randomly accessible sequence of bytes. A regular file contains binary or text data and has no structure imposed by the system. Contrast with special file.

relative pathname. In the Open System Services (OSS) file system and Network File System (NFS), a pathname that does not begin with a slash (`/`) character. A relative pathname is resolved beginning with the current working directory. Contrast with absolute pathname.

release version update (RVU). A collection of compatible revisions of HP NonStop Kernel operating system software products, identified by an RVU ID, and shipped and supported as a unit. An RVU consists of the object modules, supporting files, and documentation for the product revisions. An RVU also includes a set of documentation for the overall RVU.

RELOAD. An HP Tandem Advanced Command Language (TACL) command to load the HP NonStop™ Kernel operating system image from disk over the ServerNet fabrics into the memory of the processor.

remote access. A form of remote support, configured in OSM and TSM Notification Director. Remote access, or dial-in, allows a service provider to dial in to your system console and access your HP NonStop™ S-series server to diagnose hardware and software problems. See also remote notification.

remote interprocessor communication (RIPC). The exchange of messages between processors in different systems or nodes.

remote mount. A mount used by a Network File System (NFS) client to attach part of the local NFS file hierarchy to a point within the client's remote file hierarchy. The remote mount is visible only to the NFS client performing the mount. In effect, the local hierarchy from the mount point down is exported to the client performing the remote mount.

remote node. See remote system.

remote notification. A form of remote support. Remote notification, or dial-out, allows the OSM or TSM Notification Director to notify a service provider, such as the Global Customer Support Center (GCSC), of pending hardware and software problems. See also remote access.

remote operator. The person who performs routine system operations from a geographical distance, usually when no local operator is present.

remote procedure. A procedure or function packaged to be called within a server process indirectly by a client process.

remote procedure call. A remote procedure or the action of calling a remote procedure.

Remote Procedure Call (RPC). A protocol that extends a procedure-call form of process-to-process communication to a network environment. RPC is a way for programs running on client computers to invoke the services of a program running on a server computer. RPC allows a program to call a procedure that does not exist on the client computer.

remote procedure call system. A set of facilities that includes a programming library, network resource mapping, and binding services to provide a mechanism for a client process to execute a procedure on a remote server. A remote procedure call system is a subset of the Distributed Computing Environment (DCE) and of other products.

remote processor. A processor in a node other than the node running the ServerNet cluster monitor process (SNETMON) that is reporting status about the processor.

remote switch. An HP NonStop Cluster Switch (model 6770) or HP NonStop ServerNet Switch (model 6780) in a ServerNet cluster that is not directly connected to the server that you are logged on to. The OSM Service Connection and the TSM Service Application cannot perform any actions on a remote switch. To perform actions or get additional information on a remote switch, use the OSM Service Connection (for either the 6770 or 6780 switch) or TSM Service Application (for the 6770 switch only) to log on to a server that is directly connected to the switch.

remote system. An active ServerNet node to which the local system has active external ServerNet paths. Contrast with local system.

request packet. A ServerNet packet sent from one ServerNet device to another, requesting either a read action or a write action on the part of the receiving device. In the case of a write request, the packet contains the data to be written. The receiving device is expected to take the appropriate action and return a response packet to the device that sent the request packet. See also response packet, ServerNet packet, and ServerNet transaction.

reserved symbol. An identifier that is reserved for use by system or compiler language implementors.

resistance. The measure of opposition to current that limits the amount of current that can be produced by an applied voltage. Conductors have very little resistance; insulators have a large amount of resistance. Resistance is measured in ohms.

resource. A component of a computer system that works with other components to process transactions. Terminals, workstations, processors, memory, disk drives, processes, files, and applications are examples of resources.

response. The information or confirmation supplied by a subsystem in reaction to a command. A response is typically conveyed as one or more interprocess messages from a subsystem to an application.

response packet. A ServerNet packet returned from one ServerNet device to another, responding to an earlier received read request or write request. In the case of responding to a read request, the response packet contains the data that the requesting device wanted to have read. See also request packet, ServerNet packet, and ServerNet transaction.

RESTORE. A utility for HP NonStop™ servers that copies files from a backup tape to disk. See also BACKUP.

RFC. Request for Comments. Documents compiled by number by the Internet Engineering Task Force (IETF) that define standards for intercommunication.

RFI. See radio frequency interference (RFI).

RIPC. See remote interprocessor communication (RIPC).

RISC. See reduced instruction-set computing (RISC).

RISC processor. An instruction processing unit (IPU) that is based on reduced instruction-set computing (RISC) architecture. TNS/R processors contain RISC processors.

rld library. A library that loads position-independent code (PIC) programs and their associated dynamic-link libraries (DLLs). The `rld` library also provides the `dlopen()`, `dlclose()`, `dlresultcode()`, `dlsys()`, and `dlerror()` functions.

rms. See root mean square (rms).

robot. A media-changer device that transfers a tape cartridge to a tape drive for use and then returns the cartridge to the storage cell.

root. See root fileset and root directory. See also super ID.

root directory. In the Open System Services (OSS) file system and Network File System (NFS), a directory associated with a process that the system uses for pathname resolution when a pathname begins with a slash (`/`) character.

root fileset. For the Open System Services (OSS) file system, the fileset with the device identifier of 0, normally containing the root directory. HP recommends that this fileset be named *root*.

root mean square (rms). A measurement method used to determine the direct current (DC) equivalent value for alternating voltage and current waveforms. The rms method refers to the process of sampling a waveform, squaring the samples, averaging the samples (mean value) over the period from one cycle, and then calculating the square root of the samples. In general, rms-sensing devices are more accurate than averaging meters. Measurements from averaging meters can be as low as 30 percent of the actual current for loads with high crest factors.

root user. See super ID.

router. See ServerNet router.

router 1. See ServerNet router 1.

router 2. See ServerNet router 2.

RPC. See Remote Procedure Call (RPC).

RS-232. An industry standard for serial data transmission. It describes pin assignments, signal functions, and electrical characteristics. The current standard specifies a 25-pin connector.

RS-449. An industry standard for serial data transmission. It specifies pin assignments, signal functions, electrical characteristics, and a 37-pin connector with an optional 9-pin connector for a secondary channel.

run-time data unit (RTDU). The region of a TNS object file used to store SQL/MP source and object code. It contains embedded SQL information for clients of SQL/MP. Source RTDUs are created when a program file using embedded SQL/MP is initially compiled and linked. Object RTDUs are added to the program when the file is SQL-compiled.

run-time linker. See linker.

run-time loader. See loader.

RVU. See release version update (RVU).

S700 server. See HP NonStop™ S700 server.

S7000 server. See HP NonStop™ S7000 server.

S7400 server. See HP NonStop™ S7400 server.

S7600 server. See HP NonStop™ S7600 server.

S7800 server. See HP NonStop™ S7800 server.

S70000 server. See HP NonStop™ Sxx000 server.

S72000 server. See HP NonStop™ Sxx000 server.

S74000 server. See HP NonStop™ Sxx000 server.

S76000 server. See HP NonStop™ Sxx000 server.

S86000 server. See HP NonStop™ Sxx000 server.

Sxx000 server. See HP NonStop™ Sxx000 server.

S-series servers. See HP NonStop™ S-series servers.

SAC. See ServerNet addressable controller (SAC).

sag. A reduction in voltage, usually lasting from one cycle to a few seconds. Sags are typically caused by fault clearing or by heavy load startups.

SAN. System area network. The preferred term is fabrics.

SANMAN. See external system area network manager process (SANMAN).

saveabend file. A file containing dump information needed by the system debugging tool on a TNS or TNS/R system. In UNIX systems, such files are usually called core files or core dump files. A saveabend file is a special case of a save file. See also save file.

saved-set group ID. An Open System Services (OSS) process attribute that stores a group ID so that the group ID can later be used as the effective group ID of the process.

saved-set user ID. An Open System Services (OSS) process attribute that stores a user ID so that the user ID can later be used as the effective user ID of the process.

save file. A file created through the Inspect or Debug product. A save file contains enough information about a running process at a given time to restart the process at the same point in its execution. A save file contains an image of the process, data for the process, and the status of the process at the time the save file was created.

A save file can be created through an Inspect SAVE command at any time. A save file called a saveabend file can be created by the DMON debug monitor when a process's SAVEABEND attribute is set and the process terminates abnormally.

SBB. See ServerNet buffer board (SBB).

SBI. See ServerNet bus interface (SBI).

SC. See Subscriber Channel (SC).

scalar view of the user ID. A view of the HP NonStop™ Kernel user ID, normally used in the Open System Services (OSS) environment, that is the value (*group-number* * 256) + *user-number*. Also called the UID.

SCC. See serial communications controller (SCC).

SCF. See Subsystem Control Facility (SCF).

SCL. The mnemonic subsystem name for the ServerNet cluster subsystem.

SCP. See Subsystem Control Point (SCP).

SCSI. See small computer system interface (SCSI).

SCSI object type. The Subsystem Control Facility (SCF) object type for an Open SCSI device.

SCSI passthrough terminator. A bus-terminating plug connected between a cable and the external connector of a customer-replaceable unit (CRU). The SCSI passthrough terminator contains the necessary termination resistors required by the SCSI bus. See also terminator.

SCSI plug-in card (S-PIC). A plug-in card (PIC) for the 6760 ServerNet device adapter (ServerNet/DA) that uses a small computer system interface (SCSI) interface to connect devices to an HP NonStop™ S-series system. See also plug-in card (PIC) and fiber-optic plug-in card (F-PIC).

SCSI ServerNet addressable controller (S-SAC). A ServerNet addressable controller (SAC) that is contained within a small computer system interface (SCSI) plug-in card (S-PIC).

SCSI terminator. See terminator.

SE. System engineer. See service provider.

searchList. For each loadfile, a list that is constructed and used by the linker and loader to tell them which libraries to examine, and in which order, to locate symbol definitions needed by that loadfile. The linker and loader construct the loadfile's searchList in accordance with that loadfile's import control, which is set at link time by the loadfile's programmer. A loadfile's searchList is unaffected by the import control of any other loadfile.

SEB. See ServerNet expansion board (SEB).

SEB port. A connector on ServerNet expansion boards (SEBs) used for ServerNet links. An SEB has six emitter-coupled logic (ECL)-based ServerNet ports. See also MSEB port.

Security Manager Process (SMP). A component of the Safeguard subsystem that manages all changes to the subject and object databases and authenticates user logon attempts.

segment. In general, a contiguous sequence of logically related pages of virtual memory. The pages of the segment are individually swapped in and out of physical memory as needed. Within a loadable object file, one of the portions of the file that is mapped as one unit into virtual memory as the file is loaded. See also code segment and data segment.

selectable segment. A type of logical segment formerly known as an extended data segment. The data area for a selectable segment always begins with relative segment 4, and this area can be dynamically switched among several selectable segments by calls to the Guardian SEGMENT_USE_ procedure. The effect is similar to a rapid overlaying of one large data area. See also logical segment and flat segment.

SEM. See ServerNet extender module (SEM).

semaphore. A mechanism used to provide multiple processes with access to a shared data object.

semi-globalized. An import control characteristic of a loadfile that allows the loadfile first to obtain symbols from its own definitions and then to obtain others as for a globalized loadfile. See also searchList.

sensitive command. A Subsystem Control Facility (SCF) command that can be issued only by a user with super-group access, by the owner of the subsystem, or by a member of the group of the owner of the subsystem. The owner of a subsystem is the user who started that subsystem (or any user whose application ID is the same as the server ID—the result of a PROGID option that requires super-group access). Contrast with nonsensitive command.

separately derived power source. A facility wiring system where power is derived from a generator, transformer, or converter windings and there is no direct electrical connection, including a solidly connected grounded circuit conductor (neutral), to supply conductors originating in other facility wiring systems. Types of separately derived power sources include standby power generator, uninterruptible power supply (UPS), isolation transformer, and computer-room power center (CRPC).

serial communications controller (SCC). A type of communications controller. Each quad-integrated communications controller (QUICC) has four SCCs to handle the two Ethernet ports and the two wide area network (WAN) ports.

serial copper. A standard for physical connectivity in ServerNet I and ServerNet II networks that is available both in HP NonStop™ S-series servers and in Windows NT clusters. Serial copper uses serial encoding and supports 50 and 125 megabyte/second (MB/s) speeds. The maximum link distance at 125 MB/s is 25 meters.

serial copper PIC. See serial copper plug-in card (PIC).

serial copper plug-in card (PIC). A plug-in card (PIC) for the modular ServerNet expansion board (MSEB) and I/O multifunction (IOMF) 2 customer-replaceable unit (CRU) that supports the serial copper interface. See also serial copper and plug-in card (PIC).

serial maintenance bus (SMB). A bus that connects service processors (SPs) within an enclosure to each other and to the customer-replaceable units (CRUs) in the group.

serial maintenance bus (SMB) domain. The set of enclosures, modules, field-replaceable units (FRUs), and customer-replaceable units (CRUs) connected by a common serial maintenance bus (SMB).

server. (1) An implementation of a system used as a stand-alone system or as a node in an Expand network. (2) A combination of hardware and software designed to provide services in response to requests received from clients across a network. For example, HP NonStop™ servers provide transaction processing, database access, and other services. (3) A process or program that provides services to a client or a requester. Servers are designed to receive request messages from clients or requesters; perform the desired operations, such as database inquiries or updates, security verifications, numerical calculations, or data routing to other computer systems; and return reply messages to the clients or requesters. A server process is a running instance of a server program.

server application. An application that provides a service to a client application. An application that provides local execution of remote procedure calls is an example of a server application.

ServerNet. A communications protocol developed by HP that is used in HP NonStop™ S-series servers. See also ServerNet I and ServerNet II.

ServerNet I. The first-generation ServerNet network. ServerNet I architecture is used in current HP NonStop™ S-series servers and other products. It features 50 megabytes/second speed, 6-port ServerNet routers, 8b/9b encoding, and a 64-byte maximum packet size. See also ServerNet II.

ServerNet II. The second-generation ServerNet network. ServerNet II architecture is backward-compatible with ServerNet I architecture, and it features 125 (or 50) megabytes/second speed, 12-port ServerNet routers, 8b/9b and 8b/10b (serializer ready) encoding, and a 512-byte maximum packet size. See also ServerNet I.

ServerNet adapter. A component that connects peripheral devices to the rest of the system through a ServerNet bus interface (SBI). A ServerNet adapter is similar in function to an I/O controller logic board (LB) and backplane interconnect card (BIC) in HP NonStop™ K-series servers.

ServerNet address. A virtual memory address that, when translated to a physical address, indicates where the memory access needed by a ServerNet transaction begins. In some cases, the translation can point to some entity other than memory, such as a register. The ServerNet address is included in all ServerNet read request and write request packets.

ServerNet addressable controller (SAC). An I/O controller that is uniquely addressable by a ServerNet ID in the ServerNet fabrics. A SAC is typically implemented on some portion of a processor multifunction (PMF) customer-replaceable unit (CRU), an I/O multifunction (IOMF) CRU, or a ServerNet adapter.

ServerNet buffer board (SBB). The board that provides the ServerNet connection to and from the I/O multifunction (IOMF) customer-replaceable unit (CRU). This board replaces the processor and memory board (PMB) in the IOMF CRU.

ServerNet bus interface (SBI). The I/O control and expansion packetizer application-specific integrated circuit (ASIC).

ServerNet cable. A cable that provides ServerNet links between system enclosures.

ServerNet cluster. A network of servers (nodes) connected together using the ServerNet protocol for interprocessor communication across a cluster and within its nodes. A ServerNet cluster offers linear system expansion beyond the 8-processor or 16-processor limits of a single server, achieving comparable speeds for internal and external ServerNet communication. See also cluster and HP NonStop™ ServerNet Cluster (ServerNet Cluster).

ServerNet Cluster. See HP NonStop™ ServerNet Cluster (ServerNet Cluster).

ServerNet cluster monitor process (SNETMON). A process pair with the process name $ZZSCL that manages the state of the ServerNet cluster subsystem. Each node (system) in a ServerNet cluster must have one SNETMON process pair running.

ServerNet cluster services. The functions necessary to allow a node to join, participate in, or leave an HP NonStop™ ServerNet Cluster. These functions include monitoring and control of the physical connections to the cluster, discovery of other nodes in the cluster, and automatic recovery of failed connections.

ServerNet cluster subsystem. The subsystem managed by the ServerNet cluster monitor process (SNETMON). The subsystem name is SCL. The subsystem number is 218. The subsystem identifier is ZSCL.

ServerNet/DA. See ServerNet device adapter (ServerNet/DA).

ServerNet device. Interface logic that is associated with a specific hardware unit, such as a processor or I/O adapter, and that provides the interface to the ServerNet communications network. The responsibilities of the ServerNet device are to transform message data into ServerNet packets, to transmit those packets, to receive ServerNet packets, and to unpack the data on behalf of the associated hardware unit. See also ServerNet subdevice.

ServerNet device adapter (ServerNet/DA). A ServerNet adapter that controls external devices. The 6760 ServerNet/DA contains up to four ServerNet addressable controllers (SACs), each of which can control either disk drives or tape drives.

ServerNet device ID. See ServerNet ID.

ServerNet diagram. A graphical layout of the logical connections between objects in the system. These objects can include processors, ServerNet routers, ServerNet adapters, disks on the SCSI bus, and so on.

ServerNet dump. To copy the memory of a processor to disk by using the ServerNet fabrics.

ServerNet end device. See ServerNet device.

ServerNet expansion board (SEB). (1) A connector board that plugs in to the backplane to allow one or more ServerNet cables to exit the rear of the enclosure. The SEBs and ServerNet cables allow processors in one group to communicate with processors in another group. Each SEB provides either the ServerNet X fabric or the ServerNet Y fabric for a group. (2) A collective term for both SEBs and modular SEBs (MSEBs) when a distinction between the two types of SEBs is not required.

ServerNet extender module (SEM). Equipment that increases the distance that ServerNet signals can be transmitted over fiber-optic cables to 40 kilometers. If multimode fiber-optic (MMF) and single-mode fiber-optic (SMF) ServerNet cables are used in the same system, the SEM converts MMF signals so that they can be transmitted by SMF ServerNet cables and converts SMF signals so that they can be transmitted by MMF ServerNet cables.

ServerNet/FX adapter. A ServerNet adapter that logically extends the ServerNet X and Y fabrics to other clusters in a Fiber Optic Extension (FOX) ring by using fiber-optic lines. Two 6740 ServerNet/FX adapters are used, one for the X ring and one for the Y ring.

ServerNet/FX 2 adapter. A ServerNet adapter that logically extends the ServerNet X and Y fabrics to other clusters in a Fiber Optic Extension (FOX) ring by using fiber-optic lines. Two 6742 ServerNet/FX 2 adapters are used, one for the X ring and one for the Y ring.

ServerNet ID. A unique identifier for an addressable unit on a ServerNet communications network. A unit can have multiple ServerNet node IDs. This ID is used for routing. Each packet has a source ServerNet node ID and a destination ServerNet node ID. A pair of processors operating in duplex mode share one ServerNet node ID.

ServerNet LAN Systems Access (SLSA) subsystem. A subsystem of the HP NonStop™ Kernel operating system for configuration and management of ServerNet local area network (LAN) objects in G-series release version updates (RVUs).

ServerNet link. Two unidirectional point-to-point communication paths, one in each direction, connecting a router to a ServerNet node or another router. Each ServerNet link contains a transmit channel and a receive channel.

ServerNet memory interface (SMI). An application-specific integrated circuit (ASIC) that provides the interface between the microprocessor and the two ServerNet fabrics and main memory.

ServerNet node. A system in a ServerNet cluster. See also node.

ServerNet node number. A number that identifies a member system in a ServerNet cluster. The ServerNet node number is a simplified expression of the 6-bit node-routing ID that determines the node to which a ServerNet packet is routed. The ServerNet node number is assigned based on the port to which the node is connected on the cluster switch. The ServerNet node number, which can be viewed using the Subsystem Control Facility (SCF), the OSM Service Connection, or the TSM Service Application, is unique for each node in a ServerNet cluster.

ServerNet node routing ID. A bit field used to route ServerNet packets across the external ServerNet X and Y fabrics. The ServerNet node routing ID occupies the upper six bits of the 20-bit ServerNet ID, and it is unique for each member, or node, in a ServerNet cluster. This term is the fully qualified form of *node routing ID*.

SERVERNET object type. In the Kernel subsystem, the object type for either the $ZSNET ServerNet subsystem manager process or the ServerNet X fabric or Y fabric.

ServerNet packet. The unit of transmission in a ServerNet communications network. A ServerNet packet consists of a header, a variable-size data field, and a 32-bit cyclic redundancy check (CRC) checksum covering the entire packet. The header contains fields for control, virtual memory address, and destination and source fields to identify the processor or I/O controller transmitting and receiving the packet. See also request packet and response packet.

ServerNet port. A connector used for ServerNet links. Six ServerNet ports are located on a ServerNet expansion board (SEB). Ten ServerNet ports are located on a modular ServerNet expansion board (MSEB).

ServerNet router. An application-specific integrated circuit (ASIC) responsible for routing ServerNet packets along ServerNet links in the ServerNet fabrics, using routing information that is present within the packets. A ServerNet router acts as a fully duplex crossbar switch, able to switch any of its input ports to any of its output ports. A ServerNet router in an HP NonStop™ S-series server has either six router ports (see ServerNet router 1) or twelve router ports (see ServerNet router 2).

ServerNet router 1. A model of ServerNet router that, in an HP NonStop™ S-series server, has a total of six input and six output ports. See also ServerNet router 2.

ServerNet router 2. A model of ServerNet router that, in an HP NonStop™ S-series server, has a total of twelve input and twelve output ports. See also ServerNet router 1.

ServerNet subdevice. An I/O device that sends and receives its information through a controlling device that acts as the ServerNet device for routing purposes.

ServerNet subdevice ID. The low-order (least significant) bits of a ServerNet ID, used by a ServerNet device to distribute incoming and outgoing information between itself and its associated subdevices.

ServerNet switch. A point-to-point networking device that connects ServerNet nodes to a single fabric (X or Y) of the ServerNet communications network. The ServerNet switch routes ServerNet packets among these nodes.

ServerNet II Switch. A 12-port network switch that provides the physical junction point to enable an HP NonStop™ S-series server to connect to a ServerNet cluster. The ServerNet II Switch is a component of the HP NonStop™ Cluster Switch (model 6770).

ServerNet switch board. Component in an IOAM. The ServerNet switch board provides ServerNet connections to the I/O adapters in the IOAM. It also provides maintenance functions for the modules in which they are installed.

ServerNet transaction. The bidirectional, successful transmission of a pair of ServerNet packets between two ServerNet devices. The device that originates the transaction sends a request packet, and the device that receives the request returns a response packet. See also request packet and response packet.

ServerNet wide area network (SWAN) concentrator. (1) An HP data communications peripheral that provides connectivity to an HP NonStop™ S-series server. The SWAN concentrator supports both synchronous and asynchronous data over RS-232, RS-449, X.21, and V.35 electrical and physical interfaces. (2) A collective term for both SWAN concentrators and SWAN 2 concentrators when a distinction between the two is not required.

ServerNet wide area network (SWAN) 2 concentrator. An HP data communications peripheral that provides connectivity to an HP NonStop™ S-series server. The SWAN 2 concentrator supports both synchronous and asynchronous data over RS-232, RS-449, X.21, and V.35 electrical and physical interfaces. The SWAN 2 concentrator is the next-generation SWAN concentrator and has 12 WAN ports.

service connection. A connection between the Compaq TSM client software running on a system console and the TSM server software running on an HP NonStop™ S-series server. A service connection can be used to communicate with the server only when the HP NonStop™ Kernel operating system is running. A service connection provides a comprehensive service and maintenance picture of the server and is used to perform most service management tasks. See also low-level link.

service equipment. The necessary equipment, usually consisting of circuit breakers and their accessories, that is located near the entrance point of supply conductors. This equipment constitutes the main control and cutoff means of the supply.

service processor (SP). A physical component of the processor multifunction (PMF) customer-replaceable unit (CRU) or I/O multifunction (IOMF) CRU that controls environmental and maintenance functions (including system load functions) in the enclosure. SPs operate in pairs to provide fault tolerance. The two SPs in group 01 are designated the master service processors (MSPs). Other pairs of SPs within a system are called expansion service processors (ESPs). See also expansion service processor (ESP) and master service processor (MSP).

service provider. (1) A person trained and qualified to service field-replaceable units (FRUs). (2) An organization, such as the Global Customer Support Center (GCSC), that helps you resolve problems with your HP NonStop™ S-series server. OSM and TSM allow you to use the help of a service provider by configuring client software to support remote notification and remote access.

service side. The side of a system enclosure that contains, behind an optional door, processor multifunction (PMF) customer-replaceable units (CRUs) or I/O multifunction (IOMF) CRUs, ServerNet expansion boards (SEBs), modular SEBs (MSEBs), and ServerNet adapters. The service side is opposite the appearance side. Cables are accessed from the service side. System enclosures are typically arranged so that the service side is the least visible side.

session. In the Open System Services (OSS) environment, a set of process groups associated for job control purposes. A session can have a controlling terminal.

session leader. In the Open System Services (OSS) environment, the process that created a session.

session lifetime. In the Open System Services (OSS) environment, the period that begins when a session is created and ends when the lifetime of the last remaining process group of the session ends.

setup configuration. A simple stand-alone network used to configure the OSM or TSM environment. The setup configuration consists of the server, the primary system console, an Ethernet hub, and two local area network (LAN) cables. One cable connects the primary system console to the hub, and another cable connects the hub to a processor multifunction (PMF) customer-replaceable unit (CRU) in group 01 of the server. When OSM or TSM configuration is complete, the setup configuration can serve as the working network configuration. Variations of the setup configuration can be constructed using additional cables and hubs to increase fault tolerance.

set-user-ID program. In the Open System Services (OSS) environment, a program file that has the `S_ISUID` bit set in its file mode.

shared memory. An interprocess communication mechanism that allows two or more processes to share a given region of memory.

Shared Millicode Library. An intrinsic library containing privileged or TNS-derived millicode routines used by many native-compiled programs and by emulated TNS programs. This library includes efficient string-move operations, TNS floating-point emulation, and various privileged-only operations. These routines are mode independent. They comply with native calling conventions but can be directly invoked from any mode without changing execution modes.

shared run-time library (SRL). A collection of procedures whose code and data can be loaded and executed only at a specific assigned virtual memory address (the same address in all processes). SRLs use direct addressing and do not have run-time resolution of links to and from the main program and other independent libraries. Contrast with dynamic-link library (DLL). See also TNS shared run-time library (TNS SRL) and TNS/R native shared run-time library (TNS/R native SRL).

shell. In the Open System Services (OSS) environment, a program that interprets sequences of text input as commands. A shell can operate on an input stream, or it can interactively prompt and read commands from a terminal.

shielded twisted pair (STP). A transmission medium consisting of two twisted conductors with a foil or braid shield. Contrast with unshielded twisted pair (UTP).

shutdown file. A file invoked by the local operator or by another shutdown file that contains commands to shut down system devices, communications lines, and system and application software. Contrast with startup file.

SID. See system image disk (SID) or source ServerNet ID (SID).

signal. The method by which an environment notifies a process of an event. Signals are used to notify a process when an error that is not related to input or output has occurred. See also Open System Services (OSS) signal.

signal delivery. The time when Open System Services (OSS) takes the action appropriate for a specific process and a specific signal.

signal generation. The time when an event occurs that causes a signal for a process.

signal handler. A function or procedure that is executed when a specific signal is delivered to a specific process.

Signaling System Number 7 (SS7). The protocol used in public networks to establish connections between switches.

signal mask. The set of signals that are currently blocked from delivery to a specific process.

signal reference grid. A series of conductors, constructed of pure or composite metals (for example, copper) with good surface conductivity. A superior signal reference grid is installed on the subfloor of a computer room and connected to the raised-floor structure to establish constant and equal potential for all equipment in the computer room that is connected to it.

Simple Network Management Protocol (SNMP). An asynchronous request-response protocol used for network management. SNMP originated as a means for managing Transmission Control Protocol/Internet Protocol (TCP/IP) and Ethernet networks. OSM and TSM packages can include an SNMP-compliant interface for communication between the system console and HP NonStop™ S-series server.

single-high ServerNet adapter. A ServerNet adapter that occupies only the upper half of a ServerNet adapter slot in an HP NonStop™ S-series server. Contrast with double-high ServerNet adapter.

single-high stack. A stack that includes a base, a frame, and one system enclosure. Contrast with double-high stack.

single-mode fiber-optic (SMF) plug-in card (PIC). A plug-in card (PIC) for the modular ServerNet expansion board (MSEB) and I/O multifunction (IOMF) 2 customer-replaceable unit (CRU) that supports the single-mode fiber-optic (SMF) interface.

single-mode fiber-optic (SMF) ServerNet cable. A fiber-optic cable that allows only one mode to propagate. SMF ServerNet cable has a small-diameter core for optimized long-distance transmission. See also multimode fiber-optic (MMF) ServerNet cable.

single-wide plug-in card (PIC). A small-form-factor plug-in card (PIC) that occupies one PIC slot within a customer-replaceable unit (CRU). See also double-wide plug-in card (PIC).

sinusoidal. A waveform that can be mathematically expressed by the sine function.

SIT. See system image tape (SIT).

site update tape (SUT). One or more tapes that contain each target system's site-specific subvolume and various products. Each product contains a softdoc and a complete set of files. A SUT is delivered with every new HP NonStop™ S-series system and can be ordered whenever a new release version update (RVU) of the system software is available. A full SUT contains the current RVU of the HP NonStop Kernel operating system and all product software that has been ordered with it. A partial SUT contains a subset of products for the current RVU.

SIV. See system interrupt vector (SIV).

skin effect. The tendency of higher frequency signals to flow on the outside surface, or skin, of a conductor instead of through the entire cross-section of the conductor. The result is less total conductor area available for carrying current and an increase in the resistance of the conductor at that high signal frequency.

slot. A physical, labeled space for a customer-replaceable unit (CRU) or field-replaceable unit (FRU) that is part of a module. A module contains one or more slots.

slot location. A three-number identifier for a particular slot on a system that consists of the group number, module number, and slot number; for example, 02,01,08 (group 02, module 01, slot 08).

SLSA subsystem. See ServerNet LAN Systems Access (SLSA) subsystem.

small computer system interface (SCSI). An ANSI-standard protocol used by a controller to access a device.

SMB. See serial maintenance bus (SMB).

SMB domain. See serial maintenance bus (SMB) domain.

SMF. See storage pool.

SMF PIC. See single-mode fiber-optic (SMF) plug-in card (PIC).

SMF ServerNet cable. See single-mode fiber-optic (SMF) ServerNet cable.

SMI. See ServerNet memory interface (SMI).

SMN. The mnemonic name for the external system area network manager process (SANMAN).

SMP. See Security Manager Process (SMP).

SMT. See SWAN manager task (SMT).

snapshot. (1) A file that can be created by OSM and TSM client software to record information about the status of an HP NonStop™ S-series server, including the attributes values of all system resources, at the time it was created. The file can be forwarded to your service provider to help with troubleshooting problems. (2) For Distributed Systems Management/Software Configuration Manager (DSM/SCM), a list of the target system tape and disk locations, file fingerprints for files managed by DSM/SCM, and DSM/SCM target information. The snapshot is compiled on the target system from the target database and sent to the host system to store in the host database. An instruction to create a snapshot is part of every activation package sent from the host and can also be requested independently through the Target Interface. (3) In Visual Inspect, a disk file created by the Save Snapshot command or by DMON if the SAVEABEND attribute for a process is ON and the process abends. A snapshot is an image of the process, its data, and its status at the moment it was saved. Visual

Inspect and Inspect can be used to debug snapshots, but the Inspect product refers to a snapshot a save file.

SNDA. See ServerNet device adapter (ServerNet/DA).

SNETMON. See ServerNet cluster monitor process (SNETMON).

SNMP. See Simple Network Management Protocol (SNMP).

SNMP task. A task that runs in each ServerNet wide area network (SWAN) concentrator communications line interface processor (CLIP) as part of the WAN architecture. This task accepts and replies to Simple Network Management Protocol (SNMP) request messages.

socket. An end-point for stream-oriented communication. A socket has a file descriptor.

soft reset. An action performed on an HP NonStop™ Cluster Switch that restarts the firmware on the cluster switch but does not interfere with ServerNet passthrough data traffic.

software configuration incident report. A type of incident report that reports changes in the software configuration of the server. A software configuration incident report includes the server's software configuration file.

software product revision (SPR). The method of releasing incremental software updates on HP NonStop™ S-series systems. An SPR can include one or more corrections to code, or it can contain code that adds new function to a software product.

source ServerNet ID (SID). A field in the ServerNet packet header indicating the source of the packet.

SP. See service processor (SP).

special character. A character entered from a terminal that has an effect other than being part of the input stream from that terminal.

special file. A file in the Open System Services (OSS) file system that is not a regular file. Special files include directories, FIFOs, and character special files such as terminal device files. Contrast with regular file.

SP event message. See service processor (SP).

SPI. See Subsystem Programmatic Interface (SPI).

S-PIC. See SCSI plug-in card (S-PIC).

split-star topology. A network topology that uses up to two HP NonStop™ Cluster Switches for each external fabric. External routing is implemented between the two starred halves of a ServerNet cluster. (A starred half consists of up to eight nodes attached to one set of cluster switches.) The starred segments are joined by four-lane links. Introduced with the G06.12 release version update (RVU), the split-star topology supports up to 16 nodes. See also star topology, tri-star topology, and layered topology.

spooler. The collection of files and processes that manages the printers and print jobs on the system.

SPR. See software product revision (SPR).

SP Tool Application. A PC-based software application that you can use to request information from the master service processors (MSPs) in an HP NonStop™ S-series server. This application is intended for use only by trained service providers.

SRL. See shared run-time library (SRL).

SRM. See system resource model (SRM).

SS7. See Signaling System Number 7 (SS7).

SS7TE PIC. See SS7TE plug-in card (PIC).

SS7TE2 PIC. See SS7TE2 plug-in card (PIC).

SS7TE plug-in card (PIC). A plug-in card (PIC) used in the 6763 Common Communication ServerNet adapter (CCSA) that supports the EIA-232, EIA-449, V.35, and X.21 interfaces.

SS7TE2 plug-in card (PIC). A plug-in card (PIC) used in the 6763 Common Communication ServerNet adapter (CCSA) that supports the E1, J1, and T1 interfaces.

S-SAC. See SCSI ServerNet addressable controller (S-SAC).

SSI log. See System Service Information (SSI) log.

stackable enclosure. An enclosure that can rest on top of another enclosure. A stackable enclosure is not installed on a frame base. Contrast with base enclosure.

standby power generator. A turbine-driven or engine-driven generator that provides a backup source of power to designated loads. Often used to supplement an uninterruptible power supply (UPS) in the event of extended utility outages.

star group. One set of X and Y HP NonStop™ Cluster Switches and the ServerNet nodes (up to eight) that are connected to them. A star group can be thought of as a segment of a split-star or tri-star topology. A split-star topology can contain up to two star groups. A tri-star topology can contain up to three star groups.

start mode. An attribute of Subsystem Control Facility (SCF) PROCESS objects that controls when and if an application process starts in G-series release version updates (RVUs).

star topology. A network topology in which all nodes are connected to a central hub (HP NonStop™ Cluster Switch). Each node has its own connection to the network, so a break in the connection does not affect other nodes in the network. In a ServerNet cluster, a star topology requires one cluster switch for each external fabric and can support up to eight nodes. See also split-star topology, tri-star topology, and layered topology.

startup file. A file invoked by the local operator or by another startup file that contains commands to start up system devices, communications lines, and system and application software. Contrast with shutdown file.

state. In Subsystem Control Facility (SCF), one of the generally defined possible conditions of an object with respect to the management of that object. Examples of states are DEFINED, STARTED, and STOPPED.

static information. Information that represents the set of customer-replaceable units (CRUs) on an HP NonStop™ S-series system. Contrast with dynamic information.

static server. In the Guardian environment, a process that runs continuously and provides a specific service to other processes. A static server differs from a traditional UNIX demon in that a demon actively looks for tasks to perform. A static server performs only tasks brought to its attention by a client (requestor) process. See also demon.

STFs. See super time factors (STFs).

storage pool. A set of physical disk volumes administered as a set of logical disk volumes. A logical disk volume can span multiple physical disk volumes. When a logical disk volume becomes full, more physical disk volumes can be added.

storage-pool file. A file containing a list of disk volumes to be used by an Open System Services (OSS) fileset. As these volumes are filled, more volumes can be added to the storage-pool file.

storage subsystem. A subsystem of the HP NonStop™ Kernel operating system that handles configuration and management of disk and tape devices in G-series release version updates (RVUs).

storage subsystem manager process. The generic process that starts and manages disk and tape drives. The $ZZSTO storage subsystem manager process is started and managed by the $ZZKRN Kernel subsystem manager process through the $ZPM persistence manager process.

store and forward routing. A form of message routing whereby a router must receive an entire packet or message before it can start to forward the packet or message to the next router. Contrast with wormhole routing.

STP. See shielded twisted pair (STP).

strictly conforming POSIX.1 application. An application that requires only the facilities described in ISO/IEC IS 9945-1:1990 and the applicable computer language standards. Such an application must accept any behavior or value described in ISO/IEC IS 9945-1:1990 as unspecified or implementation-defined and, for symbolic constants, accept any value permitted by ISO/IEC IS 9945-1:1990.

structured view of the user ID. A view of the HP NonStop™ Kernel user ID, normally used in the Guardian environment, that consists of either the `group-number, user-number` pair of values or the `group-name.user-name` pair of values.

subnet. See subnetwork.

subnetwork. A physical network within an Internet protocol (IP) network. Each IP network can be divided into a number of subnetworks. Within a given network, each subnetwork is treated as a separate network. Outside the network, the subnetworks appear as part of a single network. The terms *subnetwork* and *subnet* are used interchangeably.

subnetwork address. An extension of the Internet protocol (IP) addressing scheme that allows a site to use a single IP address for multiple physical networks. A subnetwork address is created by dividing the local part of an IP address into a subnetwork number (identifying a particular subnetwork) and a host number (uniquely identifying the host system within the subnetwork). The terms *subnetwork address* and *subnet address* are used interchangeably.

SUB option. In some Subsystem Control Facility (SCF) subsystems, the designation that the object name given in a command stands not just for itself but for the names of all objects at the next-lower level in the hierarchy. The given object name can stand both for itself and for the subordinate objects, or it can stand only for the subordinate objects, depending on the value of the SUB option.

subordinate objects. In Subsystem Control Facility (SCF), objects that are logically subordinate to other objects. Some subsystems are structured hierarchically, with objects of one type logically subordinate to (that is, controlled by) an object of another type. For example, a number of subdevices can be configured on a single line. Some SCF commands include a SUB option that refers to subordinate objects.

Subscriber Channel (SC). A type of head on a fiber-optic cable in which the pins connect through a push-pull mating interface.

substate. Further information about the state of a device. The state and substate together provide information about the current condition of a device or path to a device.

SUBSYS object type. The Subsystem Control Facility (SCF) object type for most subsystems that use SCF as the user interface.

subsystem. (1) A secondary or subordinate system, usually capable of operating independently of or asynchronously with a controlling system. (2) A program or set of processes that manages a cohesive set of Subsystem Control Facility (SCF) objects.

Each subsystem has a manager through which applications can request services by issuing commands defined by that subsystem. See also subsystem manager.

Subsystem Control Facility (SCF). An interactive interface for configuring, controlling, and collecting information from a subsystem and its objects. SCF enables you to configure and reconfigure devices, processes, and some system variables while your HP NonStop™ S-series server is online.

Subsystem Control Point (SCP). The message router for Subsystem Control Facility (SCF). There can be several instances of this process. Using the Subsystem Programmatic Interface (SPI), applications send each command for a subsystem to an instance of the SCP process, which in turn sends the command to the manager process of the target subsystem. SCP also processes a few commands itself. It provides security features, version compatibility, support for tracing, and support for applications implemented as process pairs.

subsystem manager. A process that performs configuration and management functions for a subsystem.

Subsystem Programmatic Interface (SPI). A set of procedures and associated definition files and a standard message protocol used to define common message-based interfaces for communication between management applications and subsystems. It includes procedures to build and decode specially formatted messages; definition files in Transaction Application Language (TAL), COBOL85, and HP Tandem Advanced Command Language (TACL) for inclusion in programs, macros, and routines using the interface procedures; and definition files in Data Definition Language (DDL) for programmers writing their own subsystems.

subvolume. A group of related files stored on a disk. All the files have the same volume and subvolume name, but each file has a unique file identifier.

summary report. A brief informational listing of status or configuration information provided by the Subsystem Control Facility (SCF) STATUS or INFO command. Contrast with detailed report.

superblock. The part of the Open System Services (OSS) environment that contains all the information about the current state of the OSS file system. The superblock contains such items as the free list and the size of inodes.

super group. The group of user IDs that have 255 as the group number. This group has special privileges; many HP utilities have commands or functions that can be executed only by a member of the super group.

super-group user. A user who can read, write, execute, and purge most files on the system. Super-group users have user IDs that have 255 as the group number.

super ID. On HP NonStop™ systems, a privileged user who can read, write, execute, and purge all files on the system. The super ID is usually a member of a system-supervisor group.

The super ID has the set of special permissions called appropriate privileges. In the Guardian environment, the structured view of the super ID, which is (255, 255), is most commonly used. In the Open System Services (OSS) environment, the scalar view of the super ID, which is 65535, is most commonly used.

super time factors (STFs). An enhancement to the Expand product that allows the extension of automatically calculated time factors to line speeds greater than 224 kilobits/second. These time factors are logarithmic based and allow specification of a much broader range of line performance.

superuser. See super ID.

supplementary group ID. An Open System Services (OSS) process attribute that is used to determine the file-access permissions for the process.

support planner. The person who creates the operational environment for the system and is responsible for the support of the system. This person creates the startup and shutdown files, performs replacement operations, and prepares the system for upgrades and additions.

surge. An increase in the amplitude of source voltage of short duration.

SUT. See site update tape (SUT).

SVID. The *System V Interface Definition* for UNIX, published by AT&T.

SVR4. System V Release 4, a specific implementation of UNIX. See also System V.

SWAN concentrator. See ServerNet wide area network (SWAN) concentrator.

SWAN 2 concentrator. See ServerNet wide area network (SWAN) 2 concentrator.

SWAN manager task (SMT). A manager task that is provided as part of the wide area network (WAN) subsystem. The SMT runs in each communications line interface processor (CLIP) and provides a variety of management functions such as coordinating data link control (DLC) and diagnostic task downloads.

SWID. The software identification tool invoked by the SYSGENR program that audits file identification information about your software.

switch. See HP NonStop™ Cluster Switch (model 6770) and HP NonStop™ ServerNet Switch (model 6780).

switch enclosure. See cluster switch enclosure.

switch group. See cluster switch group.

switch layer. See cluster switch layer.

switch layer number. See cluster switch layer number.

switch logic board. See cluster switch logic board.

switch mode power supply. A computer power supply that uses a pulse-width modulation switching inverter and nonlinear current draw characteristics. Switch-mode power supplies are widely used because of their small size and efficiency.

switch rack. See cluster switch rack.

switch zone. See cluster switch zone.

switched connect. Connection to the Enterprise Storage System (ESS) through a Fibre Channel switch (FC switch).

symbol. (1) The symbolic name of a value, typically a function entry point or a data location. In the context of loadable libraries, symbols are defined in loadfiles and referenced in the same or other loadfiles. (2) Within the ServerNet architecture, the nine or more bits that encode 8-bit data and protocol commands.

symbolic link. In the Open System Services (OSS) file system and Network File System (NFS), a type of special file that acts as a name pointer to another file. A symbolic link contains a pathname and can be used to point to a file in another fileset. Symbolic links are not included in ISO/IEC IS 9945-1:1990. Contrast with hard link.

symbolic reference. An occurrence in code or data of the value of a symbol. The symbolic reference is bound (resolved and made usable) by assigning to it the value of a definition of that symbol. The symbol value is normally the address of a function or variable named by the symbol. In position-independent code (PIC) loadfiles, symbolic references occur only in data.

symbols region. See Inspect region.

SYSGENR. The system generation program that generates a customized version of the HP NonStop™ Kernel operating system.

SYS*nn* subvolume. A subvolume on the $SYSTEM volume where the new version of the HP NonStop™ Kernel operating system image is located. Also located on the SYS*nn* subvolume is system-dependent and release version update (RVU)-dependent software. *nn* is an octal number in the range %00 through %77.

SYSPOOL. The system data space that remains in memory after all system data structures are built by the HP NonStop™ Kernel operating system at the time of a processor load.

system. All the processors, controllers, firmware, peripheral devices, software, and related components that are directly connected together to form an entity that is managed by one HP NonStop™ Kernel operating system image and operated as one computer. See also node.

system area network management process. See external system area network manager process (SANMAN).

System Code (SC). See TNS code space.

system code. A logically distinct part of the HP NonStop™ Kernel operating system that consists of operating-system procedures shared by all processors.

system configuration database. The database file on the $SYSTEM.ZSYSCONF subvolume that contains configuration information for all system objects that can be configured by the Subsystem Control Facility (SCF). Configuration information for all system objects that can be configured during system generation is contained on the $SYSTEM.SYS*nn* subvolume. See also configuration file and SYSnn subvolume.

system console. An HP-approved personal computer used to run maintenance and diagnostic software for HP NonStop S-series servers. New system consoles are preconfigured with the required HP and third-party software. When upgrading to the latest RVU, software upgrades can be installed from the HP NonStop System Console Installer CD. System consoles communicate with NonStop S-series servers over a dedicated local area network (LAN) or a nondedicated (public) LAN. System consoles configured as the primary and backup dial-out points are referred to as the primary and backup system consoles, respectively.

system enclosure. An enclosure for system components. Processor enclosures and I/O enclosures are both system enclosures. Contrast with peripheral enclosure.

system engineer (SE). See service provider.

system entry point table (SEP table). A table used on TNS systems that stores the XEP entry value for each TNS operating system procedure entry point.

system expansion. The process of making a target system larger by adding enclosures to it. The enclosures being added can be either new enclosures or enclosures from a donor system. Contrast with system reduction.

system generation. The process of creating an operating system to support a particular system configuration and release version update (RVU).

system image disk (SID). A disk copy of the HP NonStop™ Kernel operating system produced during system configuration. The DISKGEN utility copies operating system files to the SID.

system image tape (SIT). A tape that can be used to perform a system load on a system if the system subvolume has become corrupted on both $SYSTEM disks. The tape contains a minimum set of software necessary to bring up and run the system. Use the SIT only for disaster recovery. It is not needed for normal system load. Contact the Global Customer Support Center (GCSC) before loading the system from a SIT; many additional steps are required to restore your system to working order. See also tape load.

system interrupt vector (SIV). An HP NonStop™ Kernel operating system data structure that contains the addresses of interrupt handlers, parameters passed to interrupt handlers by special interrupt microcode, and other interrupt processing information.

System Library (SL). See TNS code space.

system library. A logically distinct part of the HP NonStop™ Kernel operating system that consists of user-callable library procedures and kernel procedures.

system load. (1) To start the system; to load the HP NonStop™ Kernel operating system image into the memory of a processor. See RELOAD. (2) The process of loading the operating system. A system load changes a system from an inactive to an active (or operational) state by loading software that establishes communication between the operating system and configured system peripherals.

system-managed process. Another name for a generic process.

system manager. See manager.

system number. See Expand node number.

system operator. See local operator and remote operator.

system planner. The person who plans for the hardware and software installation of a new system or for changes to a system already installed. This person arranges for site preparation, schedules the installation, and completes the Installation Document Packet.

system process. (1) A privileged process that comes into existence at system-load time and exists continuously for a given configuration for as long as the processor remains operable. (2) An HP NonStop™ Kernel operating system process, such as the memory manager, the monitor, and the input/output (I/O) control processes. The files containing system processes are invoked by ALLPROCESSORS paragraph entries. (3) A part of a single copy of the HP NonStop Kernel operating system with Open System Services (OSS) interfaces. A system process does not have an OSS process ID.

system reduction. The process of making a donor system smaller by removing enclosures from it. The enclosures removed from a donor system might be added to a target system. Contrast with system expansion.

system resizing. See system expansion or system reduction.

system resource model (SRM). A collection of C++ objects that model the diagnostic and serviceability state behavior of the system resources discovered and managed by the Compaq TSM package. The SRM has these attributes:

Generic process name	$ZZKRN.#TSM-SRM
Process name	$ZTSM
Program file name	$SYSTEM.SYS*nn*.SRM

system serial number. A unique identifier, typically five or six alphanumeric characters, assigned to an HP NonStop™ S-series server when it is built.

System Service Information (SSI) log. An Event Management Service (EMS) log that includes information about customer-replaceable unit (CRU) removal, CRU insertion, firmware loading, security authentication, incident report dial-out authorization, incident report dial-out failure, and incident report confirmation. The SSI log can be viewed using the OSM or TSM Event Viewer.

system terminal. See system console.

System V. A version of UNIX developed and marketed originally by AT&T.

TACL. See HP Tandem Advanced Command Language (TACL).

tape bootstrap. The program on a system image tape (SIT) that reads the rest of the SIT during tape load and writes it to the system disk.

tape drive. A device that moves magnetic tape past magnetic read/write heads, which read data from or write data to the tape.

tape dump. To copy the memory of a processor to tape.

tape library. A storage device consisting of magnetic tape drives, multiple storage locations for magnetic tape cartridges, an automatic mechanism for loading the cartridges into and unloading them from the drives, and a means for an operator to load cartridges into or remove cartridges from the tape library.

tape load. A system load. A tape load is the process of reading a system image tape (SIT) and writing it to the system disk. Performing a tape load from a SIT to restore the system image files to the $SYSTEM disk is generally not recommended. Perform a tape load only with the advice of the Global Customer Support Center (GCSC) or your service provider. Loading from a tape reinitializes the disk directory.

TAPE object type. The Subsystem Control Facility (SCF) object type for all tape drives attached to your system.

target system. The computer system you make larger by adding enclosures, using a process known as system expansion. See also donor system.

TB. See terabyte (TB).

TCP. See Transmission Control Protocol (TCP).

TCP/IP. See Transmission Control Protocol/Internet Protocol (TCP/IP).

TEMPLI. The Event Management Service (EMS) template installation program that merges template object files from specified subsystems and produces resident and nonresident template files.

terabyte (TB). A unit of measurement equal to 1,099,511,627,776 bytes (1024 gigabytes). See also gigabyte (GB), kilobyte (KB), and megabyte (MB).

terminal. A type of Open System Services (OSS) character special file that conforms to the interface description in Clause 7 of ISO/IEC IS 9945-1:1990.

terminator. A resistor connected to a signal wire in a bus or network for the purpose of impedance matching to prevent reflections. SCSI chains, Ethernet cables, and some LocalTalk wiring configurations require terminators.

Tetra 8 topology. A tetrahedral topology of HP NonStop™ S-series servers that allows a maximum of four processor enclosures (eight processors) and eight I/O enclosures. Contrast with Tetra 16 topology.

Tetra 16 topology. A tetrahedral topology of HP NonStop™ S-series servers that allows a maximum of eight processor enclosures (16 processors). The maximum number of I/O enclosures allowed by the Tetra 16 topology varies depending on the release version update (RVU) and the server model. Contrast with Tetra 8 topology.

tetrahedral topology. A topology of HP NonStop™ S-series servers in which the ServerNet connections between the processor enclosures form a tetrahedron. See also tetrahedron and topology.

tetrahedron. A solid bounded by four triangular faces. In ServerNet context, a tetrahedron is four processors interconnected by ServerNet links so as to form a conceptual tetrahedron. Each processor therefore has a direct connection to the other three processors. See tetrahedral topology.

text string. A variable-length sequence of ASCII characters, defined in the CONFTEXT file, that an identifier represents. When Distributed Systems Management/Software Configuration Manager (DSM/SCM) encounters an identifier, it substitutes the associated text string for the identifier.

TF. See time factor (TF).

TFDS. See HP Tandem Failure Data System (TFDS).

TFTP. See Trivial File Transfer Protocol (TFTP).

THD. See total harmonic distortion (THD).

three point fall of potential measurement method. The measurement of a grounding electrode (such as a ground rod) where ground resistance is measured with respect to two other points. The ratio of the measurements determines the resistance of the grounding electrode.

TIM. See HP NonStop™ Technical Library (NTL).

time factor (TF). A number assigned to a line, path, or route to indicate its efficiency in transporting data. The lower the time factor, the more efficient the line, path, or route. See also surge.

TLB. See translation lookaside buffer (TLB).

TMF. See HP NonStop™ Transaction Management Facility (TMF).

TNS. HP computers that support the HP NonStop™ Kernel operating system and that are based on complex instruction-set computing (CISC) technology. TNS processors implement the TNS instruction set. Contrast with TNS/R.

TNS code segment index. A value in the range 0 through 31 that indexes a code segment within the current user code, user library, system code, or system library space. This value can be encoded in five bits.

TNS code space. One of four addressable collections of TNS object code in a TNS process. They are User Code (UC), User Library (UL), System Code (SC), and System Library (SL). UC and UL exist on a per-process basis. SC and SL exist on a per-node basis.

TNS instructions. Stack-oriented, 16-bit machine instructions defined as part of the TNS environment. On TNS systems, TNS instructions are implemented by microcode. On TNS/R systems, TNS instructions are implemented by millicode routines or by translation to an equivalent sequence of RISC instructions.

TNS library. A single, optional, TNS-compiled loadfile associated with one or more application loadfiles. If a user library has its own global or static variables, it is called a TNS shared run-time library (TNS SRL). Otherwise it is called a User Library (UL).

TNS loading. A task performed at process startup time when executing a TNS object file. This task involves mapping the TNS instructions, procedure entry point (PEP) table, and external entry point (XEP) table from a TNS object file into memory.

TNS mode. The operational environment in which TNS instructions execute by inline interpretation. See also accelerated mode and TNS/R native mode.

TNS object code. The TNS instructions that result from processing program source code with a TNS language compiler. TNS object code executes on both TNS and TNS/R systems.

TNS object file. The object file created by a TNS compiler. The file contains TNS instructions and other information needed to construct the code spaces and the initial data for a TNS process.

TNS process. A process initiated by executing a TNS or accelerated object file. A TNS process, whether accelerated or not, uses TNS register and stack conventions. Contrast with TNS/R native process.

TNS shared run-time library (TNS SRL). A shared run-time library (SRL) available to TNS processes in the Open System Services (OSS) environment. A TNS process can have only one TNS SRL. A TNS SRL is implemented as a special user library that allows shared global data.

TNS signal. A signal model available to TNS processes in the Guardian environment.

TNS stack segment. See TNS user data segment.

TNS system library. A collection of HP-supplied TNS-compiled routines available to all TNS processes. There is no per-program or per-process customization of this library. All routines are immediately available to a new process. No dynamic loading of code or creation of instance data segments is involved. See also native system library.

TNS user data segment. In a TNS process, the segment at virtual address zero. Its length is limited to 128 kilobytes. A TNS program's global variables, stack, and 16-bit heap must fit within the first 64 kilobytes. See also compiler extended-data segment.

TNS user library. A user library available to TNS processes in the Guardian environment.

TNS/R. HP computers that support the HP NonStop™ Kernel operating system and that are based on reduced instruction-set computing (RISC) technology. TNS/R processors implement the RISC instruction set and are upwardly compatible with the TNS system-level architecture. Systems with these processors include most of the HP NonStop™ servers. Contrast with TNS.

TNS/R library. A TNS/R native-mode library. For a PIC-compiled application, TNS/R libraries can be dynamic-link libraries (DLLs) or hybridized native shared runtime libraries (SRLs). For an application that is not PIC compiled, TNS/R libraries can only be native SRLs.

TNS/R native mode. The operational environment in which native-compiled RISC instructions execute. See also accelerated mode and TNS mode.

TNS/R native process. A process initiated by executing code that has been compiled directly to RISC instructions rather than to TNS instructions. Such a process uses RISC register and stack conventions and executes in TNS/R native mode.

TNS/R native shared run-time library (TNS/R native SRL). A shared run-time library (SRL) available to TNS/R native processes in both the Guardian and Open System

Services (OSS) environments. TNS/R native SRLs can be either public or private. A TNS/R native process can have multiple public SRLs but only one private SRL.

Token-Ring ServerNet adapter (TRSA). A ServerNet adapter that provides a single line from an HP NonStop™ S-series server to a token-ring network, allowing the server to act as a station on the ring. The 3862 TRSA can be configured to support network speeds of 4 megabits/second (Mbps) or 16 Mbps, and the media can be either shielded twisted pair (STP) or unshielded twisted pair (UTP).

topology. The physical layout of components that define a local area network (LAN), wide area network (WAN), or ServerNet fabric. See also star topology and tetrahedral topology.

topology branch. A processor enclosure and the I/O enclosures attached to it.

total harmonic distortion (THD). The ratio, expressed in percent, of the root mean square (rms) value for all harmonics present in the output of a power source to the total rms voltage at the output, for a pure sine-wave output. The lower the THD, the better the power source.

Total Information Manager (TIM). See HP NonStop™ Technical Library (NTL).

transformer. Equipment used to step up or step down alternating-current (AC) voltage to meet the specific requirements of the load. A transformer also provides isolation and noise-attenuation properties.

transient. A short-duration, high-amplitude impulse that is imposed on the normal voltage or current.

translation lookaside buffer (TLB). A special-purpose cache, part of the RISC processor chip, that is used in quickly translating virtual addresses to physical addresses. This rapid translation is accomplished by remembering and reusing the translations of recently referenced pages.

Transmission Control Protocol (TCP). A connection-oriented protocol that provides for the reliable exchange of data between a sending and a receiving system. TCP implements functions corresponding to the Open Systems Interconnection (OSI) reference model Layer 4, the transport layer.

Transmission Control Protocol/Internet Protocol (TCP/IP). A set of layered communication protocols for connecting workstations and larger systems in both local area networks (LANs) and wide area networks (WANs). See also HP NonStop™ TCP/IP and Parallel Library TCP/IP.

tri-star topology. A network topology that uses up to three HP NonStop™ Cluster Switches for each external fabric. External routing is implemented between the three star groups of a ServerNet cluster. (A star group consists of the eight nodes attached to one set of cluster switches.) The star groups are joined by two-lane links. Introduced with the

G06.14 release version update (RVU), the tri-star topology supports up to 24 nodes. See also split-star topology, star topology, and layered topology.

Trivial File Transfer Protocol (TFTP). A protocol defined by Request for Comments (RFC) 1350. TFTP is used as a data link control (DLC) and diagnostic task.

TRSA. See Token-Ring ServerNet adapter (TRSA).

TSM. See Compaq TSM.

TSM client software. See Compaq TSM client software.

TSM Event Viewer. See Compaq TSM Event Viewer.

TSM Low-Level Link. See Compaq TSM Low-Level Link.

TSM Notification Director. A component of the Compaq TSM client software. The TSM Notification Director receives incident reports from an HP NonStop™ S-series server, displays them, and allows you to take action or forward the incident reports to your service provider for resolution. The TSM Notification Director can be configured to run on a system console at all times, even when other TSM applications are not being used.

TSM package. See Compaq TSM package.

TSM server software. See Compaq TSM server software.

TSM Service Application. See Compaq TSM Service Application.

two-lane link. The two single-mode fiber-optic (SMF) ServerNet cables that connect the HP NonStop™ Cluster Switches on the same external fabric (for example, X1, X2, and X3) in a tri-star topology.

UCME. See uncorrectable memory error (UCME).

UID. A nonnegative integer that uniquely identifies a user within a node.

In the Open System Services (OSS) environment, the UID is the scalar view of the HP NonStop™ Kernel user ID. The UID is used in the OSS environment for functions normally associated with a UNIX user ID.

unattended site. A computer environment where no operator resides on site and the only access is from a central monitoring station.

uncorrectable memory error (UCME). An error caused by incorrect data at a particular memory location. The cause of the error is such that the error is not automatically corrected by the system, and memory replacement is required. Contrast with correctable memory error (CME).

undefined. Pertaining to the use of an incorrect value for data or the incorrect behavior of a program for which the ISO/IEC IS 9945-1:1990 standard imposes no portability requirements.

undervoltage. A negative change in the amplitude of the voltage.

unicode. A 2-octet (2-byte) character code designed to represent more alphabetic and graphic characters than allowed by the ASCII character set. When the first octet of a unicode character is zero, the second octet maps to the ISO 8859-1 (ASCII) character set. Unicode is the standard character set for encoding characters in implementations of the Java language and is the default character code set used by several Microsoft Windows operating systems.

uninterruptible power supply (UPS). The equipment used to provide an uninterruptible source of power to connected equipment if a main power outage occurs. The basic components of any UPS system are a rectifier/charger that converts alternating-current (AC) power to direct-current (DC) power, batteries that store the DC power, and an inverter that converts the DC power back into AC power for distribution to the load.

unitary segment. See segment.

unmount. To make a fileset inaccessible to the users of a node.

unplanned outage. Time during which a computer system is not capable of doing useful work because of an unplanned interruption. Unplanned interruptions can include failures caused by faulty hardware, operator error, or disaster.

unshielded twisted pair (UTP). A transmission medium consisting of two twisted conductors with no cable shielding. Contrast with shielded twisted pair (STP).

unspecified. Pertaining to the use of a correct value for data or the correct behavior of a program for which the ISO/IEC IS 9945-1:1990 standard imposes no portability requirements.

UPS. See uninterruptible power supply (UPS).

upward compatibility. The ability of a requester to operate with a server of a later revision level. In this case, the requester is upward-compatible with the server, and the server is downward-compatible with the requester. Contrast with downward compatibility.

User Code (UC). See TNS code space.

user code. A logically distinct part of the HP NonStop™ Kernel operating system that consists of the code for user processes.

user database. A database within an HP NonStop™ node that contains the user name, user ID, group ID, initial working directory, and initial user program for each user of the node.

user ID. The unique identification of a user within a node.

In the Guardian environment, the term *user ID* usually means the structured view of the HP NonStop™ Kernel user ID. In the Open System Services (OSS) environment, the term *user ID* usually means the scalar view of the HP NonStop™ Kernel user ID—a number called the UID.

User Library (UL). See TNS code space.

user library. A logically distinct part of the HP NonStop™ Kernel operating system that consists of procedures that the operating system can link to a program file at run time.

user name. A string that uniquely identifies a user within the user database for a node.

UTC. See Coordinated Universal Time (UTC).

UTP. See unshielded twisted pair (UTP).

V. See volt (V).

V.35. The International Telecommunications Union, Telecommunication Standardization Sector (ITU-T) standard for data transmission at 48 kilobits/second over 60 - 108 kilohertz group band circuits. It contains the 34-pin V.34 connector specifications normally implemented on a modular RJ-45 connector. V.35 is the equivalent of Electronics Industry Association (EIA) RS-422/RS-449.

V AC. Volts of alternating current.

vertical tetrahedron. A topology of HP NonStop™ S-series servers in which the ServerNet connections between the layers of a cluster switch group form a tetrahedron. See also tetrahedron, tetrahedral topology, and cluster switch layer.

virtual file system. In UNIX and Open System Services (OSS), a file system that allows files of a fileset to be distributed across several physical devices.

volt (V). The standard unit of measure of the potential difference that is required to move an electric charge.

volume. A logical disk, which can be one or two magnetic disk drives. In HP NonStop™ S-series systems, volumes have names that begin with a dollar sign ($), such as $DATA. See also mirrored disk or volume.

WAN. See wide area network (WAN).

WANBoot process. A process provided as part of the wide area network (WAN) subsystem that implements the BOOTP protocol and provides management functions to the WAN subsystem and the WAN products.

WAN concentrator. See ServerNet wide area network (SWAN) concentrator and ServerNet wide area network (SWAN) 2 concentrator.

WAN shared driver. A driver, provided as part of the wide area network (WAN) subsystem, that provides a simplified interface to HP NonStop™ TCP/IP for use by I/O processes. The shared driver interface is similar to that provided by DOIOPLEASE on earlier systems.

WAN subsystem. See wide area network (WAN) subsystem.

WAN subsystem manager process. A process named $ZZWAN provided as part of the wide area network (WAN) subsystem that starts and manages the WAN subsystem objects, the WAN product process, and device objects. Subsystem Control Facility (SCF) commands are directed to the WAN subsystem manager process for configuring and managing the WAN subsystem and the ServerNet wide area network (SWAN) concentrator.

WAN Wizard Pro. A graphical user interface (GUI) that guides you step-by-step through the configuration of wide area network (WAN) and local area network (LAN) software and hardware.

wide area network (WAN). A network that operates over a larger geographical area than a local area network (LAN); typically, an area with a radius greater than one kilometer. The elements of a WAN can be separated by distances great enough to require telephone communications.

wide area network (WAN) subsystem. The Subsystem Control Facility (SCF) subsystem for configuration and management of WAN objects in G-series release version updates (RVUs).

wild-card character. A character that stands for any possible character or characters in a search string or in a name applying to multiple objects. In Subsystem Control Facility (SCF) object-name templates, two wild-card characters can appear: ? (question mark) for a single character and * (asterisk) for zero or more consecutive characters. See also object-name template.

Windows Internet Name Service (WINS). A name resolution service that resolves Windows NT networking computer names to Internet protocol (IP) addresses in a routed environment. A WINS server handles name registration, queries, and release version updates (RVUs). See also IP address.

WINS. See Windows Internet Name Service (WINS).

work files. Temporary files created during system generation that serve as storage areas. Work files are useful for debugging purposes after system generation. You can choose to make specified work files permanent.

working directory. In the Open System Services (OSS) environment, a directory, associated with a process, that is used in pathname resolution for relative pathnames.

worldwide name (WWN). A unique, 64-bit number assigned to hardware ports.

WORM. See write once, read many (WORM).

wormhole routing. A technique for reducing network latency in a router. Packet bytes are immediately switched to the appropriate output port as soon as they arrive rather than accumulated in a buffer until an entire packet is received. Contrast with store and forward routing.

write once, read many (WORM). A media storage class in which data, once written, cannot be erased or overwritten.

WWN. See worldwide name (WWN).

WWN zones. Similar to Ethernet virtual LANs (VLANs), WWN zones organize the cabling and interactions between components in a storage and server system.

wye. A polyphase electrical supply where the conductors of the source transformer are connected to the terminals in a physical arrangement that resembles the letter *Y*. Each point of the Y represents the connection for a conductor at high potential. The angle of phase displacement between each point on the Y is 120 degrees. The center point of the Y is the common return point for the neutral conductor.

wye-delta. The interconnections between a wye source and a delta load.

X.21. A digital signaling interface recommended by the International Telecommunications Union, Telecommunication Standardization Sector (ITU-T) that includes specifications for data terminal equipment/data communications equipment (DTE/DCE) physical interface elements, alignment of call control characters and error checking, elements of the call control phase for circuit switched services, data transfer at or below 9600 bits/second, and test loops.

X fabric. The X side of the internal or external ServerNet fabrics. See also fabric, external ServerNet X or Y fabric, and internal ServerNet X or Y fabric.

XIO. See extensible input/output (XIO).

XLLINK. The linker program invoked during system generation to link accelerated (file code 100) TNS object files to create system code and system library files.

XO bond. The bond connection on an isolation transformer, installed between the transformer's neutral XO terminal and ground. This bond is required at North American and British sites to provide an effective return path for any equipment conductor fault current back to the neutral of the isolation transformer.

Y fabric. The Y side of the internal or external ServerNet fabrics. See also fabric, external ServerNet X or Y fabric, and internal ServerNet X or Y fabric.

$YMIOP. The name of the maintenance I/O process that is built during system generation and is available at system startup.

$ZCNF. The name of the configuration utility process.

zero-signal reference. A connecting point, bus, or conductor used as one side of a signal circuit. Such a reference object might or might not necessarily be designated as a ground. Sometimes referred to as a common circuit.

$ZEXP. The name of the Expand manager process.

$ZM*nn*. The name of the QIO monitor process in processor *nn*.

$ZNET. The name of the Subsystem Control Point (SCP) management process.

zombie process. In the Open System Services (OSS) environment, a process that has terminated but is still recorded in system tables.

zone. See cluster switch zone and WWN zones.

$ZPM. The name of the persistence manager process.

ZSCL. The subsystem identifier for the ServerNet cluster subsystem.

ZSERVER. The object file name of the $ZSVR server process for the labeled-tape subsystem.

ZSMN. The subsystem identifier for the external system area network manager process (SANMAN).

$ZSVR. The name of the server process for the labeled-tape subsystem. See also ZSERVER.

ZSYSCONF subvolume. The subvolume on the $SYSTEM disk that contains the system configuration database.

$ZTC0. The default transport-provider process that provides Transmission Control Protocol/Internet Protocol (TCP/IP) services to Open System Services (OSS) `AF_INET` sockets programs.

$ZZATM. The name of the Asynchronous Transfer Mode (ATM) monitor process.

$ZZFOX. The name of the Fiber Optic Extension (FOX) monitor process in the ServerNet/FX adapter subsystem.

$ZZKRN. The name of the Kernel subsystem manager process.

$ZZLAN. The name of the ServerNet LAN Systems Access (SLSA) subsystem manager process that is started by the $ZZKRN Kernel subsystem manager process and maintained by the $ZPM persistence manager process. See also LAN manager (LANMAN) process.

$ZZPAM. The name of the Port Access Method (PAM) manager process.

$ZZSCL. The name of the ServerNet cluster monitor process (SNETMON).

$ZZSMN. The name of the external system area network manager process (SANMAN).

$ZZSTO. The name of the storage subsystem manager process.

$ZZWAN. The name of the wide area network (WAN) subsystem manager process.

Index

E

F

G

I

K

L

M

N

O

R

S

W

Z