

HP Integrity NonStop NS-Series Operations Guide

Abstract

This guide describes how to perform routine system hardware operations for HP Integrity NonStop™ NS-series servers. These tasks include monitoring the system, performing common operations tasks, and performing routine hardware maintenance. This guide is written for system operators.

Product Version

N.A.

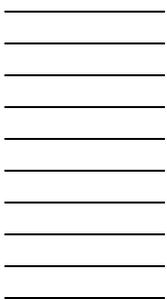
Supported Release Version Updates (RVUs)

This guide supports H06.08 and all subsequent H-series RVUs until otherwise indicated by its replacement publication.

Part Number	Published
529869-005	November 2006

Document History

Part Number	Product Version	Published
529869-003	N.A.	February 2006
529869-004	N.A.	August 2006
529869-005	N.A.	November 2006



HP Integrity NonStop NS-Series Operations Guide

Index	Examples	Figures	Tables
-----------------------	--------------------------	-------------------------	------------------------

- [What's New in This Manual](#) xiii
 - [Manual Information](#) xiii
 - [New and Changed Information](#) xiii
- [About This Guide](#) xv
 - [Who Should Use This Guide](#) xv
 - [What Is in This Guide](#) xvi
 - [Where to Get More Information](#) xvii
 - [Notation Conventions](#) xviii

1. Introduction to Integrity NonStop NS-Series Operations

- [When to Use This Section](#) 1-2
- [Understanding the Operational Environment](#) 1-2
- [What Are the Operator Tasks?](#) 1-2
 - [Monitoring the System and Performing Recovery Operations](#) 1-2
 - [Preparing for and Recovering from Power Failures](#) 1-3
 - [Stopping and Powering Off the System](#) 1-3
 - [Powering On and Starting the System](#) 1-3
 - [Creating Startup and Shutdown Files](#) 1-3
 - [Performing Preventive Maintenance](#) 1-3
 - [Operating Disk Drives and Tape Drives](#) 1-3
 - [Responding to Spooler Problems](#) 1-4
 - [Updating Firmware](#) 1-4
- [Determining the Cause of a Problem:](#)
 - [A Systematic Approach](#) 1-4
 - [A Problem-Solving Worksheet](#) 1-4
 - [Task 1: Get the Facts](#) 1-6
 - [Task 2: Find and Eliminate the Cause of the Problem](#) 1-7
 - [Task 3: Escalate the Problem If Necessary](#) 1-8
 - [Task 4: Prevent Future Problems](#) 1-9

Logging On to an Integrity NonStop Server	1-9
System Consoles	1-9
Opening a TACL Window	1-10
Overview of OSM Applications	1-11
Launching OSM Applications	1-11
Service Procedures	1-12
Support and Service Library	1-12

2. Determining Your System Configuration

When to Use This Section	2-1
Modular Hardware Components	2-2
Differences Between Integrity NonStop NS-Series Systems	2-2
Terms Used to Describe System Hardware Components	2-4
Recording Your System Configuration	2-4
Using SCF to Determine Your System Configuration	2-5
SCF System Naming Conventions	2-5
SCF Configuration Files	2-5
Using SCF to Display Subsystem Configuration Information	2-6
Displaying SCF Configuration Information for Subsystems	2-9
Additional Subsystems Controlled by SCF	2-13
Displaying Configuration Information—SCF Examples	2-15

3. Overview of Monitoring and Recovery

When to Use This Section	3-1
Functions of Monitoring	3-2
Monitoring Tasks	3-2
Working With a Daily Checklist	3-2
Tools for Checking the Status of System Hardware	3-3
Additional Monitoring Tasks	3-6
Monitoring and Resolving Problems—An Approach	3-7
Using OSM to Monitor the System	3-7
Using the OSM Service Connection	3-7
Recovery Operations for Problems Detected by OSM	3-12
Monitoring Problem Incident Reports	3-12
Using SCF to Monitor the System	3-12
Determining Device States	3-13
Automating Routine System Monitoring	3-16
Using the Status LEDs to Monitor the System	3-20
Related Reading	3-22

4. Monitoring EMS Event Messages

- [When to Use This Section](#) 4-1
- [What Is the Event Management Service \(EMS\)?](#) 4-1
- [Tools for Monitoring EMS Event Messages](#) 4-1
 - [OSM Event Viewer](#) 4-2
 - [EMSDIST](#) 4-2
 - [ViewPoint](#) 4-2
 - [Web ViewPoint](#) 4-2
- [Related Reading](#) 4-2

5. Processes: Monitoring and Recovery

- [When to Use This Section](#) 5-1
- [Types of Processes](#) 5-1
 - [System Processes](#) 5-1
 - [I/O Processes \(IOPs\)](#) 5-2
 - [Generic Processes](#) 5-2
- [Monitoring Processes](#) 5-3
 - [Monitoring System Processes](#) 5-3
 - [Monitoring IOPs](#) 5-4
 - [Monitoring Generic Processes](#) 5-4
- [Recovery Operations for Processes](#) 5-6
- [Related Reading](#) 5-6

6. Communications Subsystems: Monitoring and Recovery

- [When to Use This Section](#) 6-1
- [Communications Subsystems](#) 6-1
 - [Local Area Networks \(LANs\) and Wide Area Networks \(WANs\)](#) 6-2
- [Monitoring Communications Subsystems and Their Objects](#) 6-4
 - [Monitoring the SLSA Subsystem](#) 6-4
 - [Monitoring the WAN Subsystem](#) 6-6
 - [Monitoring the NonStop TCP/IP Subsystem](#) 6-9
 - [Monitoring Line-Handler Process Status](#) 6-10
 - [Tracing a Communications Line](#) 6-12
- [Recovery Operations for Communications Subsystems](#) 6-13
- [Related Reading](#) 6-13

7. ServerNet Resources: Monitoring and Recovery

- [When to Use This Section](#) 7-1
- [ServerNet Communications Network](#) 7-1
- [System I/O ServerNet Connections](#) 7-4
- [Monitoring the Status of the ServerNet Fabrics](#) 7-4
 - [Monitoring the ServerNet Fabrics Using OSM](#) 7-5
 - [Monitoring the ServerNet Fabrics Using SCF](#) 7-6
- [Related Reading](#) 7-8

8. I/O Adapters and Modules: Monitoring and Recovery

- [When to Use This Section](#) 8-1
- [I/O Adapters and Modules](#) 8-2
 - [Fibre Channel ServerNet Adapter \(FCSA\)](#) 8-2
 - [Gigabit Ethernet 4-Port Adapter \(G4SA\)](#) 8-2
 - [4-Port ServerNet Extender \(4PSE\)](#) 8-3
- [Monitoring I/O Adapters and Modules](#) 8-3
 - [Monitoring the FCSAs](#) 8-4
 - [Monitoring the G4SAs](#) 8-5
 - [Monitoring the 4PSEs](#) 8-7
- [Recovery Operations for I/O Adapters and Modules](#) 8-7
- [Related Reading](#) 8-8

9. Processors and Components: Monitoring and Recovery

- [When to Use This Section](#) 9-1
- [Overview of the NonStop Blade Complex](#) 9-2
- [Monitoring and Maintaining Processors](#) 9-4
 - [Monitoring Processors Automatically Using TFDS](#) 9-4
 - [Monitoring Processor Status Using the OSM Low-Level Link](#) 9-5
 - [Monitoring Processor Status Using the OSM Service Connection](#) 9-5
 - [Monitoring Processor Performance Using ViewSys](#) 9-7
- [Identifying Processor Problems](#) 9-7
 - [Processor or System Hangs](#) 9-7
 - [Processor Halts](#) 9-8
 - [OSM Alarms and Attribute Values](#) 9-8

Recovery Operations for Processors	9-9
Recovery Operations for a Processor Halt	9-9
Halting One or More Processors	9-10
Reloading a Single Processor on a Running Server	9-10
Recovery Operations for a System Hang	9-14
Enabling/Disabling Processor and System Freeze	9-15
Freezing the System and Freeze-Enabled Processors	9-15
Dumping a Processor to Disk	9-15
Backing Up a Processor Dump to Tape	9-19
Replacing Processor Memory	9-19
Replacing the Processor Board and Processor Entity	9-19
Submitting Information to Your Service Provider	9-19
Related Reading	9-22

[10. Disk Drives: Monitoring and Recovery](#)

When to Use This Section	10-1
Overview of Disk Drives	10-2
Internal SCSI Disk Drives	10-2
M8xxx Fibre Channel Disk Drives	10-3
Enterprise Storage System (ESS) Disks	10-3
Monitoring Disk Drives	10-4
Monitoring Disk Drives With OSM	10-4
Monitoring Disk Drives With SCF	10-5
Monitoring the State of Disk Drives	10-9
Monitoring the Use of Space on a Disk Volume	10-9
Monitoring the Size of Database Files	10-9
Monitoring Disk Configuration and Performance	10-10
Identifying Disk Drive Problems	10-11
Internal SCSI Disk Drives	10-11
M8xxx Fibre Channel Disk Drives	10-11
Recovery Operations for Disk Drives	10-12
Recovery Operations for a Down Disk or Down Disk Path	10-14
Recovery Operations for a Nearly Full Database File	10-15
Related Reading	10-15

11. Tape Drives: Monitoring and Recovery

- [When to Use This Section](#) 11-1
- [Overview of Tape Drives](#) 11-1
- [Monitoring Tape Drives](#) 11-2
 - [Monitoring Tape Drive Status With OSM](#) 11-2
 - [Monitoring Tape Drive Status With SCF](#) 11-5
 - [Monitoring Tape Drive Status With MEDIACOM](#) 11-6
 - [Monitoring the Status of Labeled-Tape Operations](#) 11-7
- [Identifying Tape Drive Problems](#) 11-7
- [Recovery Operations for Tape Drives](#) 11-8
 - [Recovery Operations Using the OSM Service Connection](#) 11-8
 - [Recovery Operations Using SCF](#) 11-9
- [Related Reading](#) 11-9

12. Printers and Terminals: Monitoring and Recovery

- [When to Use This Section](#) 12-1
- [Overview of Printers and Terminals](#) 12-1
- [Monitoring Printer and Collector Process Status](#) 12-2
 - [Monitoring Printer Status](#) 12-2
 - [Monitoring Collector Process Status](#) 12-2
- [Recovery Operations for Printers and Terminals](#) 12-3
 - [Recovery Operations for a Full Collector Process](#) 12-3
- [Related Reading](#) 12-3

13. Applications: Monitoring and Recovery

- [When to Use This Section](#) 13-1
- [Monitoring TMF](#) 13-1
 - [Monitoring the Status of TMF](#) 13-2
 - [Monitoring Data Volumes](#) 13-2
 - [TMF States](#) 13-3
- [Monitoring the Status of Pathway](#) 13-4
 - [PATHMON States](#) 13-5
- [Related Reading](#) 13-6

14. Power Failures: Preparation and Recovery

When to Use This Section	14-2
System Response to Power Failures	14-2
NonStop NS-Series Cabinets (Modular Cabinets)	14-2
NonStop S-Series I/O Enclosures	14-2
External Devices	14-2
ESS Cabinets	14-3
Air Conditioning	14-3
Preparing for Power Failure	14-3
Set Ride-Through Time	14-3
Configure OSM Power Fail Support	14-3
Monitor Power Supplies	14-4
Monitor Batteries	14-4
Maintain Batteries	14-4
Power Failure Recovery	14-4
Procedure to Recover From a Power Failure	14-5
Setting System Time	14-5
Related Reading	14-5

15. Starting and Stopping the System

When to Use This Section	15-2
Powering On a System	15-2
Powering On the System From a Low Power State	15-3
Powering On the System From a No Power State	15-3
Starting a System	15-5
Loading the System	15-5
Starting Other System Components	15-9
Performing a System Load	15-9
Performing a System Load From a Specific Processor	15-11
Reloading Processors	15-12
Minimizing the Frequency of Planned Outages	15-14
Anticipating and Planning for Change	15-14
Stopping Application, Devices, and Processes	15-14
Stopping the System	15-16
Alerts	15-16
Halting All Processors Using OSM	15-16
Powering Off a System	15-17
System Power-Off Using OSM	15-17
System Power-Off Using SCF	15-17
Emergency Power-Off Procedure	15-18

Troubleshooting and Recovery Operations	15-18
Fans Are Not Turning	15-18
System Does Not Appear to Be Powered On	15-19
Green LED Is Not Lit After POSTs Finish	15-19
Amber LED on a Component Remains Lit After the POST Finishes	15-19
Components Fail When Testing the Power	15-19
Recovering From a System Load Failure	15-20
Getting a Corrupt System Configuration File Analyzed	15-21
Recovering From a Reload Failure	15-21
Exiting the OSM Low-Level Link	15-22
Opening Startup Event Stream and Startup TACL Windows	15-22
Related Reading	15-24

16. Creating Startup and Shutdown Files

Automating System Startup and Shutdown	16-2
Managed Configuration Services (MCS)	16-2
Startup	16-2
Shutdown	16-3
For More Information	16-3
Processes That Represent the System Console	16-3
\$YMIOP.#CLCI	16-3
\$YMIOP.#CNLS	16-3
\$ZHOME	16-4
\$ZHOME Alternative	16-4
Example Command Files	16-4
CIIN File	16-5
Establishing a CIIN File	16-6
Modifying a CIIN File	16-6
If a CIIN File Is Not Specified or Enabled in OSM	16-7
Example CIIN Files	16-8
Writing Efficient Startup and Shutdown Command Files	16-9
Command File Syntax	16-9
Avoid Manual Intervention	16-10
Use Parallel Processing	16-10
Investigate Product-Specific Techniques	16-11
How Process Persistence Affects Configuration and Startup	16-11
Tips for Startup Files	16-11

Startup File Examples	16-12
System Startup File	16-12
Spooler Warm-Start File	16-14
TMF Warm-Start File	16-14
TCP/IP Stack Configuration and Startup File	16-14
CP6100 Lines Startup File	16-17
ATP6100 Lines Startup File	16-17
X.25 Lines Startup File	16-17
Printer Line Startup File	16-18
Expand-Over-IP Line Startup File	16-18
Expand Direct-Connect Line Startup File	16-18
Tips for Shutdown Files	16-19
Shutdown File Examples	16-19
System Shutdown File	16-20
CP6100 Lines Shutdown File	16-21
ATP6100 Lines Shutdown File	16-21
X.25 Lines Shutdown File	16-21
Printer Line Shutdown File	16-22
Expand-Over-IP Line Shutdown File	16-22
Direct-Connect Line Shutdown File	16-22
Spooler Shutdown File	16-23
TMF Shutdown File	16-23

17. Preventive Maintenance

When to Use This Section	17-1
Monitoring Physical Facilities	17-1
Checking Air Temperature and Humidity	17-1
Checking Physical Security	17-2
Maintaining Order and Cleanliness	17-2
Checking Fire-Protection Systems	17-2
Cleaning System Components	17-2
Cleaning an Enclosure	17-2
Cleaning and Maintaining Printers	17-2
Cleaning Tape Drives	17-3
Handling and Storing Cartridge Tapes	17-3

A. Operational Differences Between Systems Running G-Series and H-Series RVUs

B. Tools and Utilities for Operations

When to Use This Appendix	B-1
BACKCOPY	B-2
BACKUP	B-2
Disk Compression Program (DCOM)	B-2
Disk Space Analysis Program (DSAP)	B-2
EMSDIST	B-2
Event Management Service Analyzer (EMSA)	B-2
File Utility Program (FUP)	B-3
Measure	B-3
MEDIACOM	B-3
NonStop NET/MASTER	B-3
NSKCOM and the Kernel-Managed Swap Facility (KMSF)	B-3
OSM Package	B-3
PATHCOM	B-4
PEEK	B-4
RESTORE	B-4
SPOOLCOM	B-4
Subsystem Control Facility (SCF)	B-4
HP Tandem Advanced Command Language (TACL)	B-5
TMFCOM	B-5
Web ViewPoint	B-5
ViewPoint	B-5
ViewSys	B-6

C. Related Reading

D. Converting Numbers

When to Use This Appendix	D-1
Overview of Numbering Systems	D-2
Binary to Decimal	D-3
Octal to Decimal	D-4
Hexadecimal to Decimal	D-5
Decimal to Binary	D-7
Decimal to Octal	D-8
Decimal to Hexadecimal	D-9

Safety and Compliance

Index

Examples

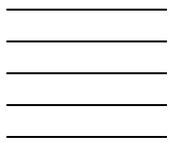
Example 2-1.	SCF LISTDEV Command Output	2-7
Example 2-2.	SCF ADD DISK Command Output	2-11
Example 2-3.	SCF INFO PROCESS Command Output	2-15
Example 2-4.	SCF INFO SAC Command Output	2-15
Example 2-5.	SCF INFO PROCESS \$ZZWAN Command Output	2-16
Example 2-6.	SCF INFO LINE Command Output	2-16
Example 3-1.	SCF STATUS TAPE Command	3-13
Example 3-2.	System Monitoring Command File	3-16
Example 3-3.	System Monitoring Output File	3-17

Figures

Figure 3-1.	OSM Management: System Icons Indicate Problems Within	3-8
Figure 3-2.	Expanding the Tree Pane to Locate the Source of Problems	3-9
Figure 3-3.	Attributes Tab	3-10
Figure 3-4.	Using System Status Icons to Monitor Multiple Systems	3-10
Figure 3-5.	Alarm Summary Dialog Box	3-11
Figure 3-6.	Problem Summary Dialog Box	3-11
Figure 7-1.	Integrity NonStop NS16000 System	7-2
Figure 7-2.	Integrity NonStop NS14000 System with IOAM Enclosure	7-3
Figure 7-3.	I/O Connections to the PICS in a P-Switch	7-4
Figure 9-1.	Modular NSAA With One NonStop Blade Complex and Four Processors	9-3
Figure 9-2.	Processor Status Display	9-5
Figure 9-3.	OSM Representation of Processor Complex	9-6
Figure 11-1.	OSM: Monitoring Tape Drives Connected to an FCSA	11-3
Figure 11-2.	OSM: Monitoring Tape Drives Connected to an IOMF2	11-4
Figure 15-1.	System Load Dialog Box	15-10
Figure 15-2.	Logical Processor Reload Parameters	15-13
Figure 15-3.	Opening a Startup TACL Window	15-22
Figure 15-4.	OutsideView Buttons on the Windows Toolbar	15-22
Figure D-1.	Binary to Decimal Conversion	D-3
Figure D-2.	Octal to Decimal Conversion	D-4
Figure D-3.	Hexadecimal to Decimal Conversion	D-6

Tables

Table 1-1.	Problem-Solving Worksheet	1-5
Table 2-1.	Key Subsystems and Their Logical Device Names and Device Types	2-8
Table 2-2.	Displaying Information for the TCP/IP Subsystem (\$ZTCO)	2-9
Table 2-3.	Displaying Information for the Kernel Subsystem (\$ZZKRN)	2-10
Table 2-4.	Displaying Information for the Storage Subsystem (\$ZZSTO)	2-10
Table 2-5.	Displaying Information for the SLSA Subsystem (\$ZZLAN)	2-12
Table 2-6.	Displaying Information for the WAN Subsystem (\$ZZWAN)	2-13
Table 2-7.	Subsystem Objects Controlled by SCF	2-13
Table 3-1.	Monitoring System Components	3-4
Table 3-2.	Daily Tasks Checklist	3-6
Table 3-3.	SCF Object States	3-14
Table 3-4.	Status LEDs and Their Functions	3-20
Table 3-5.	Related Reading for Monitoring	3-22
Table 4-1.	Related Reading for Monitoring EMS Event Messages	4-2
Table 6-1.	Related Reading for Communications Lines and Devices	6-13
Table 8-1.	Service, Flash Firmware, Flash Boot Firmware, Device, and Enabled States for the FCSA	8-4
Table 8-2.	Service, Device, and Enabled States for the G4SA	8-6
Table 8-3.	Related Reading for I/O Adapters and Modules	8-8
Table 9-1.	Other Files to Submit to Your Service Provider	9-20
Table 9-2.	Additional Processor Dump Information for Your Service Provider	9-21
Table 9-3.	Related Reading for Monitoring and Recovery Operations on Processors	9-22
Table 10-1.	Primary and Backup Path States for Disk Drives	10-9
Table 10-2.	Possible Causes of Common Disk Drive Problems	10-11
Table 10-3.	Common Recovery Operations for Disk Drives	10-12
Table 11-1.	Common Tape Drive Problems	11-7
Table 11-2.	Related Reading for Tapes and Tape Drives	11-9
Table 13-1.	TMF States	13-3
Table 15-1.	System Load Paths in Order of Use	15-7
Table 15-2.	Related Reading for Starting and Stopping a System	15-24
Table C-1.	Related Reading for Tools and Utilities	C-1
Table D-1.	Descriptions of Number Systems	D-2



What's New in This Manual

Manual Information

Abstract

This guide describes how to perform routine system hardware operations for HP Integrity NonStop™ NS-series servers. These tasks include monitoring the system, performing common operations tasks, and performing routine hardware maintenance. This guide is written for system operators.

Product Version

N.A.

Supported Release Version Updates (RVUs)

This guide supports H06.08 and all subsequent H-series RVUs until otherwise indicated by its replacement publication.

Part Number	Published
529869-005	November 2006

Document History

Part Number	Product Version	Published
529869-003	N.A.	February 2006
529869-004	N.A.	August 2006
529869-005	N.A.	November 2006

New and Changed Information

This manual has been updated to include references to HP Integrity NonStop NS14000 and NS1000 servers containing VIO enclosures (in place of an IOAM enclosure).

About This Guide

This guide describes how to perform routine system hardware operations for HP Integrity NonStop NS-series servers on H-series release version updates.

This guide is primarily geared toward commercial type NonStop NS-series servers (see [Differences Between Integrity NonStop NS-Series Systems](#) on page 2-2 for high-level architectural and hardware differences between the various commercial models). While basic monitoring principles, such as [Using OSM to Monitor the System](#) on page 3-7, apply to Telco as well as commercial systems, refer to the *NonStop NS-Series Carrier Grade Server Manual* for hardware details and service procedures specific to Telco systems.

Note. *NS-series* refers to the hardware that makes up the server. *H-series* refers to the software that runs on the server.

The term, *NonStop server*, refers to both NonStop S-series servers and Integrity NonStop NS-series servers.

Use this guide along with the *Guardian User's Guide* and the written policies and procedures of your company regarding:

- General operations
- Security
- System backups
- Starting and stopping applications

Who Should Use This Guide

This guide is written for operators who perform system hardware operations. It provides an overview of the routine tasks of monitoring the system and guides the operator through the infrequent tasks of starting and stopping the system and performing online recovery on the system.

What Is in This Guide

Section or Appendix	Section and Appendix Titles
Section 1	<u>Introduction to Integrity NonStop NS-Series Operations</u>
Section 2	<u>Determining Your System Configuration</u>
Section 3	<u>Overview of Monitoring and Recovery</u>
Section 4	<u>Monitoring EMS Event Messages</u>
Section 5	<u>Processes: Monitoring and Recovery</u>
Section 6	<u>Communications Subsystems: Monitoring and Recovery</u>
Section 7	<u>ServerNet Resources: Monitoring and Recovery</u>
Section 8	<u>I/O Adapters and Modules: Monitoring and Recovery</u>
Section 9	<u>Processors and Components: Monitoring and Recovery</u>
Section 10	<u>Disk Drives: Monitoring and Recovery</u>
Section 11	<u>Tape Drives: Monitoring and Recovery</u>
Section 12	<u>Printers and Terminals: Monitoring and Recovery</u>
Section 13	<u>Applications: Monitoring and Recovery</u>
Section 14	<u>Power Failures: Preparation and Recovery</u>
Section 15	<u>Starting and Stopping the System</u>
Section 16	<u>Creating Startup and Shutdown Files</u>
Section 17	<u>Preventive Maintenance</u>
Appendix A	<u>Operational Differences Between Systems Running G-Series and H-Series RVUs</u>
Appendix B	<u>Tools and Utilities for Operations</u>
Appendix C	<u>Related Reading</u>
Appendix D	<u>Converting Numbers</u>

Where to Get More Information

Operations planning and operations management practices appear in these manuals:

- *NonStop NSxxxx Planning Guide* for your NS16000, NS14000, or NS1000 server
- *Availability Guide for Application Design*
- *Availability Guide for Change Management*
- *Availability Guide for Problem Management*

Note. For manuals not available in the H-series collection, please refer to the G-series collection on NTL.

For comprehensive information about performing operations tasks for an Integrity NonStop NS-series server, you need both this guide and the *Guardian User's Guide*. The *Guardian User's Guide* describes some tasks not covered in this guide, such as supporting users of the system.

The *Guardian User's Guide* describes routine tasks common to system operations on all NonStop servers. Instructions and examples show how to support users of the system, how to monitor operator messages, how to control the spooler, and how to manage disks and tapes. Numerous tools that support these functions are also documented. Some monitoring procedures in the *Guardian User's Guide* have information about using only the Subsystem Control Facility (SCF). That guide does not generally describe any monitoring procedures using the OSM packages.

Information about the use of OSM, such as how to migrate from TSM to OSM, how to install and configure OSM server and client components, and how to use the OSM Service Connection, appear in these manuals:

- *OSM Migration and Configuration Guide*
- *NonStop System Console Installer Guide*
- *OSM Service Connection User's Guide* (available in NTL and as online help within the OSM Service Connection)

Servers that are connected in ServerNet clusters require special installation and operating procedures that are not documented in this manual. Such information is instead provided with the appropriate cluster documentation and the *ServerNet Cluster Supplement for Integrity NonStop NS-Series Servers*.

In the 6780 ServerNet cluster environment, installation and operating procedures are documented in these manuals:

- *ServerNet Cluster 6780 Planning and Installation Guide*
- *ServerNet Cluster 6780 Operations Guide*

Installation and operating procedures for earlier server clusters (those using 6770 switches) are documented in:

- *ServerNet Cluster Manual*

OSM is the required system management tool for servers that use 6780 switches in ServerNet clusters, but OSM also provides system management for earlier versions of ServerNet clusters.

For other documentation related to operations tasks, refer to [Appendix C, Related Reading](#).

Support and Service Library

These NTL Support and Service library categories provide procedures, part numbers, troubleshooting tips, and tools for servicing NonStop S-series and Integrity NonStop NS-series systems:

- Hardware Service and Maintenance Publications
- Service Information
- Service Procedures
- Tools and Download Files
- Troubleshooting Tips

Within these categories, where applicable, content might be further categorized according to server or enclosure type.

Authorized service providers can also order the NTL Support and Service Library 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://hps.knowledgemanagement.hp.com/wow/order.asp>.

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 DAM Volumes and 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 ]
```

A group of items enclosed in brackets is a list from which you can choose one item or none. The items in the list may 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:

```
FC [ num ]
   [ -num ]
   [ text ]
```

```
K [ X | D ] address
```

{ } Braces. A group of items enclosed in braces is a list from which you are required to choose one item. The items in the list may be arranged either vertically, with aligned braces on each side of the list, or horizontally, enclosed in a pair of braces and separated by vertical lines. For example:

```
LISTOPENS PROCESS { $appl-mgr-name }
                  { $process-name }
```

```
ALLOWSU { ON | OFF }
```

| Vertical Line. A vertical line separates alternatives in a horizontal list that is enclosed in brackets or braces. For example:

```
INSPECT { OFF | ON | SAVEABEND }
```

... **Ellipsis.** An ellipsis immediately following a pair of brackets or braces indicates that you can repeat the enclosed sequence of syntax items any number of times. For example:

```
M address [ , new-value ]...
[ - ] {0|1|2|3|4|5|6|7|8|9}...
```

An ellipsis immediately following a single syntax item indicates that you can repeat that syntax item any number of times. For example:

```
"s-char..."
```

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
```

Quotation marks around a symbol such as a bracket or brace indicate the symbol is a required character that you must enter as shown. For example:

```
"[ repetition-constant-list ]"
```

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
```

Line Spacing. If the syntax of a command is too long to fit on a single line, each continuation line is indented three spaces and is separated from the preceding line by a blank line. This spacing distinguishes items in a continuation line from items in a vertical list of selections. For example:

```
ALTER [ / OUT file-spec / ] LINE
    [ , attribute-spec ]...
```

Notation for Messages

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

Bold Text. Bold text in an example indicates user input entered at the terminal. For example:

```
ENTER RUN CODE
?123
CODE RECEIVED:      123.00
```

The user must press the Return key after typing the input.

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]

{ } Braces. A group of items enclosed in braces is a list of all possible items that can be displayed, of which one is actually displayed. The items in the list might be arranged either vertically, with aligned braces on each side of the list, or horizontally, enclosed in a pair of braces and separated by vertical lines. For example:

obj-type obj-name state changed to *state*, caused by
{ Object | Operator | Service }

process-name State changed from *old-objstate* to *objstate*
{ Operator Request. }
{ Unknown. }

| Vertical Line. A vertical line separates alternatives in a horizontal list that is enclosed in brackets or braces. For example:

Transfer status: { OK | Failed }

% Percent Sign. A percent sign precedes a number that is not in decimal notation. The % notation precedes an octal number. The %B notation precedes a binary number. The %H notation precedes a hexadecimal number. For example:

%005400

%B101111

%H2F

P=%*p-register* E=%*e-register*

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:

The message types specified in the REPORT clause are different in the COBOL85 environment and the Common Run-Time Environment (CRE).

The CRE has many new message types and some new message type codes for old message types. In the CRE, the message type SYSTEM includes all messages except LOGICAL-CLOSE and LOGICAL-OPEN.

1

Introduction to Integrity NonStop NS-Series Operations

[When to Use This Section](#) on page 1-2

[Understanding the Operational Environment](#) on page 1-2

[What Are the Operator Tasks?](#) on page 1-2

[Monitoring the System and Performing Recovery Operations](#) on page 1-2

[Preparing for and Recovering from Power Failures](#) on page 1-3

[Stopping and Powering Off the System](#) on page 1-3

[Powering On and Starting the System](#) on page 1-3

[Performing Preventive Maintenance](#) on page 1-3

[Operating Disk Drives and Tape Drives](#) on page 1-3

[Responding to Spooler Problems](#) on page 1-4

[Updating Firmware](#) on page 1-4

[Determining the Cause of a Problem: A Systematic Approach](#) on page 1-4

[A Problem-Solving Worksheet](#) on page 1-4

[Task 1: Get the Facts](#) on page 1-6

[Task 2: Find and Eliminate the Cause of the Problem](#) on page 1-7

[Task 3: Escalate the Problem If Necessary](#) on page 1-8

[Task 4: Prevent Future Problems](#) on page 1-9

[Logging On to an Integrity NonStop Server](#) on page 1-9

[System Consoles](#) on page 1-9

[Opening a TACL Window](#) on page 1-10

[Overview of OSM Applications](#) on page 1-11

[Launching OSM Applications](#) on page 1-11

[Service Procedures](#) on page 1-12

[Support and Service Library](#) on page 1-12

When to Use This Section

This section introduces system hardware operations for Integrity NonStop NS-series servers. It provides an introduction to the other sections in this guide.

Understanding the Operational Environment

To understand the operational environment:

- If you are already familiar with other NonStop systems, see [Appendix A, Operational Differences Between Systems Running G-Series and H-Series RVUs](#).
- For a brief introduction to the system organization and the location of system components in an Integrity NonStop server, see [Section 2, Determining Your System Configuration](#).
- For information about various software tools and utilities you can use to perform system operations on an Integrity NonStop server, see [Appendix B, Tools and Utilities for Operations](#).

What Are the Operator Tasks?

The system operations described in this guide include:

- Monitoring the system and performing recovery operations
- Preparing for and recovering from power failures
- Stopping and powering off the system
- Powering on and starting the system
- Performing preventive maintenance
- Operating disk drives and tape drives
- Responding to spooler problems

Monitoring the System and Performing Recovery Operations

Checking for indications of potential system problems by monitoring the system is part of the normal system operations routine. You perform recovery operations to restore a malfunctioning system component to normal use. Most recovery procedures for Integrity NonStop servers can be performed online. Monitoring the status of all system components and performing recovery operations are described in:

- [Section 3, Overview of Monitoring and Recovery](#)
- [Section 4, Monitoring EMS Event Messages](#)
- [Section 5, Processes: Monitoring and Recovery](#)
- [Section 6, Communications Subsystems: Monitoring and Recovery](#)
- [Section 7, ServerNet Resources: Monitoring and Recovery](#)
- [Section 8, I/O Adapters and Modules: Monitoring and Recovery](#)
- [Section 9, Processors and Components: Monitoring and Recovery](#)

- [Section 10, Disk Drives: Monitoring and Recovery](#)
- [Section 11, Tape Drives: Monitoring and Recovery](#)
- [Section 12, Printers and Terminals: Monitoring and Recovery](#)
- [Section 13, Applications: Monitoring and Recovery](#)

Recovery operations for a system console are not discussed in this guide. For recovery procedures for a system console and the applications installed on the system console, see the *NonStop NSxxxx Hardware Installation Manual* for your Integrity NonStop NS16000, NS14000, or NS1000 server.

Preparing for and Recovering from Power Failures

You can minimize unplanned outage time by having procedures to prepare and recover quickly from power failures, as described in [Section 14, Power Failures: Preparation and Recovery](#).

Stopping and Powering Off the System

HP recommends a specific set of procedures for stopping and powering off an Integrity NonStop server or its components, as described in [Section 15, Starting and Stopping the System](#).

Powering On and Starting the System

HP recommends a specific set of procedures for powering on and starting an Integrity NonStop server or its components, as described in [Section 15, Starting and Stopping the System](#).

Creating Startup and Shutdown Files

HP recommends a specific set of procedures for creating startup and shutdown files on an Integrity NonStop server or its components, as described in [Section 16, Creating Startup and Shutdown Files](#).

Performing Preventive Maintenance

Routine preventive maintenance consists of:

- Dusting or cleaning enclosures as needed
- Cleaning tape drives regularly
- Evaluating tape condition regularly
- Cleaning and reverifying tapes as needed

Routine hardware maintenance procedures are described in [Section 17, Preventive Maintenance](#).

Operating Disk Drives and Tape Drives

Refer to the documentation shipped with the drive.

Responding to Spooler Problems

Refer to the *Spooler Utilities Reference Manual*.

Updating Firmware

Refer to the *H06.xx Software Installation and Upgrade Guide*

Determining the Cause of a Problem: A Systematic Approach

Continuous availability of your NonStop system is important to system users, and your problem-solving processes can help make such availability a reality. To determine the cause of a problem on your system, start by trying the easiest, least expensive possibilities. Move to more complex, expensive possibilities only if the easier solutions fail.

This subsection presents an approach you can use in your operations environment to:

- Determine the possible causes of problems
- Systematically fix or escalate such problems
- Develop ways of preventing the same problems from recurring

The four basic steps in systematic problem solving are:

Task	Page
Task 1: Get the Facts	1-6
Task 2: Find and Eliminate the Cause of the Problem	1-7
Task 3: Escalate the Problem If Necessary	1-8
Task 4: Prevent Future Problems	1-9

A Problem-Solving Worksheet

[Table 1-1](#) is a worksheet that you can use to help you through the problem-solving process. Use this worksheet to:

- Get the facts about a problem
- Find and eliminate the cause of the problem
- Make any appropriate escalation decisions
- Prevent future problems

Make copies of this worksheet and use it to collect and analyze facts regarding a problem you are experiencing. The results might not tell you exactly what is occurring, but they will narrow down the number of possible causes.

You are authorized by HP to reproduce this worksheet only for the purpose of operating your system.

Table 1-1. Problem-Solving Worksheet

Problem Facts	Possible Causes
What?	
Where?	
When?	
Magnitude?	
Situation Facts	Escalation Decision
Plan to Verify/Fix	
Plan to Prevent and Control Damage	

Task 1: Get the Facts

The first step in solving any problem is to get the facts. Although it is tempting to speculate about causes, your time is better spent in first understanding the symptoms of the problem.

Task 1a: Determine the Facts About the Problem

To get a clear, complete description of problem symptoms, ask questions to determine the facts about the problem. For example:

Category	Questions to Ask
What?	What are you having trouble with? What specifically is wrong?
Where?	Where did you first notice the problem? Where has it occurred since you first noticed it? Which applications, components, devices, and people are affected?
When?	When did the problem occur? What is the frequency of the problem? Has this problem occurred before this time?
Magnitude?	Is the problem quantifiable in any way? (That is, can it be measured?) For example, how many people are affected? Is this problem getting worse?

Task 1b: Determine the Facts About the Situation

Collect facts about the situation in which the problem arose. A clear description of the situation that led to the problem could indicate a simple solution. Examples of questions to ask are:

- Who reported the problem and how can this person be contacted?
- How critical is the situation?
- What events led to the problem?
- Has anything changed recently that might have caused the problem?
- What event messages have you received?
- What is the current configuration of the hardware and software products affected?

An example of information you might obtain from asking questions:

Question	Answer
What is happening that indicates a problem?	A terminal is hung.
Where is this problem occurring?	In the office of USER.BONNIE. The affected terminal is named \$JT1.#C02.
When is this problem occurring?	At 8:30 this morning and also at the same time two days ago. Both times, this problem occurred after three unsuccessful attempts to log on.
What is the magnitude of this problem?	Intermittent; the problem seemed to disappear on its own when it first occurred two days ago.

Task 2: Find and Eliminate the Cause of the Problem

After you collect the facts, you are ready to begin considering the possible causes of a problem. Using these facts and relying on your knowledge and experience, begin to list possible causes of the problem.

Task 2a: Identify the Most Likely Cause

To evaluate the possible causes of any problem, you must compare each cause with the problem symptoms. The problem-solving worksheet gives you a guide for accomplishing this task. In the following example:

- Possible causes become column headings
- Entries made in the worksheet's rows indicate whether the cause in that column could have produced the problem symptoms you listed in that row.
 - Write *yes* in the appropriate box if that cause could explain that symptom.
 - Write *no* in the appropriate box if a possible cause does not explain a fact.

The most likely cause is the one that best explains all the facts; that is, the cause that contains the most *yes* answers.

For example, possible causes of a hung terminal problem could be:

- A terminal hardware problem
- A stopped or suspended TACL process
- System security, which locks a user out after three unsuccessful logon attempts

This worksheet lists some possible causes of a hung terminal and illustrates further how to evaluate the possible causes:

Problem Facts	Possible Causes		
	Terminal hardware	TACL process	Security
What? Terminal \$JT1.#C02 is hung	Yes	Yes	Yes
Where? Office of USER.BONNIE	Yes	Yes	Yes
When? 8:30 a.m. today Two days ago at 8:30 a.m. After 3 failed logon attempts	Yes Yes No	Yes Yes No	Yes Yes Yes
Magnitude? Intermittent Goes away on its own	? ?	Yes Yes	Yes Yes

Task 2b: Fix the Most Probable Cause of the Problem

For the example in the worksheet, the most likely cause of the hung terminal is a security problem. Ask yourself what would be the fastest, least expensive, safest, and surest way of verifying that this is the most probable cause of the problem.

Once you have determined the most likely cause, try to fix it. Follow through and implement the appropriate solution. If this solution does not fix the problem, continue trying other possible solutions that are reasonable considering time, expense, and safety.

Task 3: Escalate the Problem If Necessary

If the solutions you tried in the previous tasks do not solve the problem, you might consider escalating the problem to get additional help.

Task 3a: Determine Whether You Need to Escalate the Problem

After you complete each task in the problem-solving process, you must decide whether you can continue by yourself or if you must ask for help. Ask yourself these questions:

- Do I have the authority to resolve this problem?
- Do I have the necessary knowledge?
- Do I have the skill?
- Do I have the time?
- What other people need to become involved, if any?
- Who needs to be informed about the problem's status?

Task 3b: Provide Documentation

If you decide to escalate the problem, you might be required to document the problem by providing:

- A problem identification number
- A problem classification
- A complete description and history of the problem
- Diagnostic information such as copies of the event log, results of memory dumps, and so on

You might also have procedures at your site for logging problems. If you have a shift log or problem log, make timely entries in the log.

Task 4: Prevent Future Problems

Solving problems that occur with your system can be exciting because it is active and stimulating. Preventing problems is often less dramatic. But in the end, prevention is more productive than solving problems. The more work you do to prevent problems before they arise, the fewer problems that will arise at potentially critical times.

These questions provide a framework for your problem-prevention efforts:

- Why did this problem occur? What was the root cause? Were there any contributing causes?
- How serious was the problem?
- What is the likelihood that it will occur again?
- Is it possible to eliminate the causes of this problem?
- Is it possible to reduce the likelihood that this problem will occur in the future?
- Can automation tools be used to detect and respond to preliminary symptoms of this problem?
- Can anything be done now to minimize the damage that would result from a reoccurrence of this problem?
- Can the problem resolution process be improved in any way?

Logging On to an Integrity NonStop Server

Many operations and troubleshooting tasks are performed by logging on to your Integrity NonStop server from a system console and using the TACL command interpreter or one of the OSM applications. For example, the TACL command interpreter allows you to access SCF, which you use to configure, control, and collect information about objects within subsystems. For examples of OSM tasks and functions, see [Overview of OSM Applications](#) on page 1-11.

System Consoles

A system console is a personal computer approved by HP to run maintenance and diagnostic software for Integrity NonStop 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 Integrity NonStop servers over a dedicated service LAN (local area network). System consoles configured as the primary and backup dial-out points are referred to as the primary and backup system consoles, respectively.

The OSM Low-Level Link and OSM Notification Director applications reside on the system console, along with other required HP and third-party software. OSM Service Connection and OSM Event Viewer software resides on your server, and connectivity

is established from the console through Internet Explorer browser sessions. For more information, see [Launching OSM Applications](#) on page 1-11.

Opening a TACL Window

On a system console, you must open a TACL window before you can log on to the TACL command interpreter. For information about logging on to a TACL command interpreter, see the *Guardian User's Guide*.

You can use any of the following methods to open a TACL window.

Opening a TACL Window Directly From OutsideView

If you know the IP address of the NonStop server (not that of OSM), use this method:

1. Select **Start>Programs>OutsideView32 7.1**.
2. From the **Session** menu, select **New**. The New Session Properties dialog box appears.
3. From the New Session Properties dialog box, Session tab, click **IO Properties**. The TCP/IP Properties dialog box appears.
4. In the TCP/IP Properties dialog box:
 - a. In the Host name or IP address and port box, type the IP address, followed by a space and the port number. For example:

```
172.17.22.187 23
```

The port number is 23 for a TACL prompt and 301 for a Startup TACL prompt. In general, you should use port number 23 to perform operations tasks.
 - b. Click **OK**.
5. From the New Session Properties dialog box, click **OK**. A TACL window appears.
6. Log on to the TACL prompt.

Opening a TACL Window From the Low-Level Link

You can also open a TACL window from the OSM Low-Level Link application as described in the Troubleshooting section in [Opening Startup Event Stream and Startup TACL Windows](#) on page 15-22.

For more details on the functions of the TACL command interpreter, see [Appendix B, Tools and Utilities for Operations](#).

Overview of OSM Applications

HP NonStop Open System Management (OSM) applications perform a variety of functions, such as:

- The OSM Low-Level Link Application is primarily used for down-system support, such as [Two startup event stream windows and two startup TACL windows are automatically launched on the system console configured to receive them.](#) on page 15-6, [Recovery Operations for Processors](#) on page 9-9, and configuring IOAM, VIO, and P-switch modules (see the *NonStop NSxxxx Hardware Installation Manual* for your Integrity NonStop NS16000, NS14000, or NS1000 server).
- The OSM Service Connection is used to monitor, inventory, and perform actions on system and ServerNet Cluster components. See [Using OSM to Monitor the System](#) on page 3-7 for an overview of how the OSM Service Connection is used to monitor your system components.
- The OSM Event Viewer is used for [Section 4, Monitoring EMS Event Messages.](#)
- The OSM Notification Director is used for [Monitoring Problem Incident Reports](#) on page 3-12 and dialing out information to your service provider.

Launching OSM Applications

Several operations tasks in this guide require you to log on to one of the OSM applications. Assuming that all OSM client components have been installed on the system console, launch the desired application as described below, then see the online help (or default home page, for the browser-based OSM applications) for log-on instructions.

To launch OSM applications: **Start>Programs>HP OSM.** Then select the name of the application to launch:

- OSM Service Connection
- OSM Low-Level Link Application
- OSM Notification Director>Start/Stop
- OSM Event Viewer
- OSM System Inventory Tool

The OSM Service Connection and the OSM Event Viewer are browser-based applications. Assuming that the OSM Console Tools component has been installed on the system console, the Start menu shortcuts launch a default web page for these two applications. From that page, you can select the system of your choice from the list of bookmarks displayed in the left column of the page (available bookmarks include those that were user-created during previous sessions and those converted automatically from an existing OSM system list). If no bookmarks are available, the web page also contains instructions on how to access these applications by entering a system URL as an Internet Explorer address. The system console-based OSM Console Tools component is not required to use the OSM Service Connection and the OSM Event

Viewer applications; it merely installs the Start menu shortcuts and default home pages to make accessing these applications easier. You can also simply open a new Internet Explorer browser window and enter the URL of the system you wish to access.

For more information on configuring, accessing, or using OSM applications, see:

- *OSM Migration and Configuration Guide*
- *OSM Service Connection User's Guide*
- Online help within the OSM Service Connection, Low-Level Link, Notification Director, and Event Viewer applications

Service Procedures

OSM offers a variety of guided procedures, interactive actions, and documented service procedures to automate or assist with system serviceability. They are launched by actions within the OSM Service Connection, and include online help.

For a list (and help files) for service procedures, both those incorporated into OSM and others that are not part of OSM, refer to the Support and Service Library.

Support and Service Library

These NTL Support and Service library categories provide procedures, part numbers, troubleshooting tips, and tools for servicing NonStop S-series and Integrity NonStop NS-series systems:

- Hardware Service and Maintenance Publications
- Service Information
- Service Procedures
- Tools and Download Files
- Troubleshooting Tips

Within these categories, where applicable, content might be further categorized according to server or enclosure type.

Authorized service providers can also order the NTL Support and Service Library 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://hps.knowledgemanagement.hp.com/wow/order.asp>.

2

Determining Your System Configuration

[When to Use This Section](#) on page 2-1

[Modular Hardware Components](#) on page 2-2

[Differences Between Integrity NonStop NS-Series Systems](#) on page 2-2

[Terms Used to Describe System Hardware Components](#) on page 2-4

[Recording Your System Configuration](#) on page 2-4

[Using SCF to Determine Your System Configuration](#) on page 2-5

[SCF System Naming Conventions](#) on page 2-5

[SCF Configuration Files](#) on page 2-5

[Using SCF to Display Subsystem Configuration Information](#) on page 2-6

[Displaying SCF Configuration Information for Subsystems](#) on page 2-9

[Additional Subsystems Controlled by SCF](#) on page 2-13

[Displaying Configuration Information—SCF Examples](#) on page 2-15

When to Use This Section

This section describes the system enclosures, the system organization, numbering and labeling, and how to identify components in an Integrity NonStop NS-series server. For detailed information on system hardware organization, refer to the *NonStop NSxxxx Planning Guide* for your Integrity NonStop NS16000, NS14000, or NS1000 server.

Modular Hardware Components

Hardware for Integrity NonStop systems is implemented in modules or enclosures that are installed in modular cabinets. The servers include these hardware components:

- Modular Cabinet with Power Distribution Unit (PDU)
- NonStop Blade Complex
- NonStop Blade Element
- Logical Synchronization Unit (LSU) (in Integrity NonStop NS16000 and NS14000 systems only; Integrity NonStop NS1000 systems have no LSUs)
- Processor Switch, or p-switch (in Integrity NonStop NS16000 systems only; Integrity NonStop NS14000 and NS1000 systems have no processor switches)
- I/O Adapter Module (IOAM) Enclosure, including subcomponent I/O Adapters:
 - Fibre Channel ServerNet adapter (FCSA)
 - Gigabit Ethernet 4-port ServerNet adapter (G4SA)
 - 4-Port ServerNet Extenders (4PSEs) (Integrity NonStop NS14000 and NS1000 systems only)
- VIO Enclosure (displayed by OSM as a VIO Module object) — For more information, see [Integrity NonStop NS14000 Systems](#), [Integrity NonStop NS1000 Systems](#), or the *Versatile I/O (VIO) Manual*.
- Fibre Channel disk module (FCDM)
- Maintenance Switch (Ethernet)
- UPS and ERM
- NonStop System Console (to manage the system)
- Cable Management Devices
- Enterprise Storage System (ESS)

Differences Between Integrity NonStop NS-Series Systems

NonStop System Architectures

Integrity NonStop NS-series systems offer a variety of architecture and configuration options to suit different customer needs. Integrity NonStop NS16000 and Integrity NonStop NS14000 systems take advantage of NonStop advanced architecture (NSAA). For more information, see the *NonStop NS16000 Planning Guide* or *NonStop NS14000 Planning Guide*. Integrity NonStop NS1000 systems employ the NonStop value architecture (NSVA). For more information, see the *NonStop NS1000 Planning Guide*.

Integrity NonStop NS16000 Systems

In Integrity NonStop NS16000 systems, IOAM enclosures connect through ServerNet links to the processors via the processor switches. One IOAM enclosure provides ServerNet connectivity for up to 10 ServerNet I/O adapters on each of the two ServerNet fabrics. FCSAs and G4SAs can be installed in an IOAM enclosure for

communications to storage devices and subsystems as well as to LANs. Additional IOAM enclosures can be added to increase connectivity and storage resources.

Integrity NonStop NS16000 systems connect to NonStop S-series I/O enclosures by using fiber-optic ServerNet links to connect the p-switches of the Integrity NonStop system to IOMF2 CRUs in the I/O enclosures.

Integrity NonStop NS14000 Systems

In Integrity NonStop NS14000 systems, there are no p-switches. There are now two types of NS14000 systems:

- A NonStop NS14000 system consisting of a single IOAM enclosure, with an I/O adapter module on each ServerNet fabric — processor connections are made through ports on 4-Port ServerNet Extenders (4PSEs), located in slot one and optionally slot 2 of each I/O adapter module, to the processors via the LSUs. The IOAM enclosure provides ServerNet connectivity for up to 8 ServerNet I/O adapters on each of the two ServerNet fabrics (FCSAs and G4SAs can be installed in slots 2 through 5 of the two IOAMs in the IOAM enclosure for communications to storage devices and subsystems as well as to LANs). Integrity NonStop NS14000 systems do not support connections to additional IOAM enclosures or NonStop S-series I/O enclosures.
- A NonStop NS14000 system consisting of two VIO enclosures, one on each ServerNet fabric — processor connections for processors 0-3 are made through ports 1-4 of the VIO Logic Board in slot 14 of each VIO enclosure, via the LSUs. An optional Optical Extender PIC in slot 2 provides for additional processor connectivity (processors 4-7). VIO enclosures have embedded ports and allow for optional expansion ports to supply the equivalent functionality provided by FCSAs and G4SAs in NS14000 systems with IOAMs.

Integrity NonStop NS14000 systems do not support connections to additional IOAM enclosures or NonStop S-series I/O enclosure

For more information on Integrity NonStop NS14000 systems, see the *Versatile I/O (VIO) Manual*, the *NonStop NS14000 Planning Guide*, or the *NonStop NS14000 Hardware Installation Manual*.

Integrity NonStop NS1000 Systems

Integrity NonStop NS1000 systems have no processor switches or LSUs. Like Integrity NonStop NS14000 systems, there are now two types: those consisting of a single IOAM enclosure (two IOAMs) and those consisting of one VIO enclosure for each fabric. ServerNet connectivity for each type is accomplished as described for the [Integrity NonStop NS14000 Systems](#), except for the absence of the LSUs.

Integrity NonStop NS1000 systems do not support connections to NonStop S-series I/O enclosures. Besides the architectural differences, Integrity NonStop NS1000 systems also utilize different NonStop Blade Elements than Integrity NonStop NS16000 or NS14000 systems. For more information on Integrity NonStop NS1000 systems, refer to the *NonStop NS1000 Planning Guide* and the *NonStop NS1000 Hardware Installation Manual*.

Terms Used to Describe System Hardware Components

The terms used to describe system hardware components vary. These terms include:

- Device
- System resource or object

Device

A device can be a physical device or a logical device. A physical device is a physical component of a computer system that is used to communicate with the outside world or to acquire or store data. A logical device is a process used to conduct input or output with a physical device.

System Resource or Object

The term “system resource” is used in OSM documentation to refer to server components that OSM software displays, monitors, and often controls. The term “object” is often used when referring to a specific resource, such as “the Disk object.” All system resources are displayed in hierarchical form in the tree pane of the OSM Service Connection; many are also displayed in Physical or Inventory views of the view pane. The effect of selecting an object in either pane is the same: for example, you can view attributes for the selected system resource in the Attributes tab, view alarms for that resource (if any exist) in the Alarms tab, or right-click on the resource object and select Actions, to display the Actions dialog box (from which you can select and perform actions on the selected system resource). Besides physical hardware components, such as IOAM enclosures, power supplies, ServerNet adapters, and disk and tape drives, system resources also include logical entities that OSM supports, such as logical processors, ServerNet fabrics, and LIFs (logical interfaces).

Recording Your System Configuration

As a system operator, you need to understand how your system is configured so you can confirm that the hardware and system software are operating normally. If problems do occur, knowing your configuration allows you to pinpoint problems more easily. If your system configuration is corrupted, documentation about your configuration is essential for recovery. You should be familiar with the system organization, system configuration, and naming conventions.

Several methods are available for researching and recording your system configuration:

- Maintaining records in hard-copy format
- Using the OSM Service Connection to inventory your system
In the OSM Service Connection tree pane, select the System object. From the View pane drop-down menu, select Inventory to display a list of the system’s hardware resources. Click Save to save this list to a Microsoft Excel file.
- Using SCF to list objects and devices and to display subsystem configuration information

For information on forms available that can help you record your system configuration, refer to the *NonStop NSxxxx Planning Guide* for your Integrity NonStop NS16000, NS14000, or NS1000 server.

Using SCF to Determine Your System Configuration

SCF is one of the most important tools available to you as a system operator. SCF commands configure and control the objects (lines, controllers, processes, and so on) belonging to each subsystem running on the Integrity NonStop NS-series server. You also use SCF to display information about subsystems and their objects.

SCF accepts commands from a workstation, a disk file, or an application process. It sends display output to a workstation, a file, a process, or a printer. Some SCF commands are available only to some subsystems. An overall SCF reference is the *SCF Reference Manual for H-Series RVUs*. Subsystem-specific information appears in a separate manual for each subsystem. For a partial list of these manuals, refer to [Appendix C, Related Reading](#).

More details about the functions of SCF appear in [Subsystem Control Facility \(SCF\)](#) on [page B-4](#).

SCF System Naming Conventions

SCF object names usually follow a consistent set of naming conventions defined for each installation. HP preconfigures some of the naming conventions to create the logical device names for many SCF objects.

System planning and configuration staff at your site likely will change or expand on the preconfigured file-naming conventions that HP provides, typically by establishing naming conventions for configuring such objects as storage devices, communication processes, and adapters. These conventions should simplify your monitoring tasks by making process or object functions intuitively obvious to someone looking at the object name. For example, in your environment, tape drives might be named \$TAPE n , where n is a sequential number.

The *SCF Reference Manual for H-Series RVUs* lists naming conventions for SCF objects, as well as HP reserved names that cannot be changed or used for other objects or processes in your environment.

SCF Configuration Files

Your system is delivered with a standard set of configuration files:

- The \$SYSTEM.SYS nn .CONFBASE file contains the minimal configuration required to load the system.
- The \$SYSTEM.ZSYSCONF.CONFIG file contains a standard system configuration created by HP. This basic configuration includes such objects as disk drives, tape

drives, ServerNet adapters, the local area network (LAN) and wide area network (WAN) subsystem manager processes, the OSM server processes, and so on. You typically use this file to load the system.

- The `$SYSTEM.ZSYSCONF.CONFIG` file is also saved on your system as the `ZSYSCONF.CONF0000` file.

All subsequent changes to the system configuration are made using SCF. The system saves configuration changes on an ongoing basis in the `ZSYSCONF.CONFIG` file. You have the option to save a stable copy of your configuration at any time in `ZSYSCONF.CONFxyy` using the SCF SAVE command. For example:

```
-> SAVE CONFIGURATION 01.02
```

You can save multiple system configurations by numbering them sequentially based on a meaningful convention that reflects, for example, different hardware configurations. Each time you load the system from `CONFBASE` or `CONFxyy`, the system automatically saves in a file called `ZSYSCONF.CONFSAVE` a copy of the configuration file used for the system load.

For guidelines on how to recover if your system configuration files are corrupted, refer to [Troubleshooting and Recovery Operations](#) on page 15-18.

For certain SCF subsystems, configuration changes are persistent. The changes persist through processor and system loads unless you load the system with a different configuration file. Examples of these subsystems are the Kernel, ServerNet LAN Systems Access (SLSA), the storage subsystem, and WAN. For other SCF subsystems, the changes are not persistent. You must reimplement them after a system or processor load. Examples of these subsystems are General Device Support (GDS), Open System Services (OSS), and SQL communication subsystem (SCS).

Using SCF to Display Subsystem Configuration Information

SCF enables you to display, in varying levels of detail, the configuration of objects in each subsystem supported by SCF. For example, you can use the `LISTDEV` command to list all the devices on your system or to list the objects within a given subsystem. Then you can use the `INFO` command with a logical device name or device type to obtain information about a specific device or class of devices.

Another useful command when displaying information is the `ASSUME` command. Use the `ASSUME` command to define a current default object and fully qualified object name. Then you can use `INFO` to display information just for that object. For example, if you type this command and then enter the `INFO` command without specifying an object, SCF displays only the information for the workstation called `$LI.#TERM1`:

```
> SCF ASSUME WS $L1.#TERM1
```

SCF LISTDEV: Listing the Devices on Your System

To obtain listings for most devices and processes that have a device type known to SCF, at a TACL prompt type:

```
> SCF LISTDEV
```

In the example shown in [Example 2-1](#), the SCF LISTDEV command lists all the physical and logical devices on the system.

Example 2-1. SCF LISTDEV Command Output

```
$SYSTEM STARTUP 1> SCF LISTDEV
```

LDev	Name	PPID	BPID	Type	RSize	Pri	Program
0	\$0	0,3	1,3	(1,0)	102	201	\DRP14.\$SYSTEM.SYS00.OPCOLL
1	\$NCP	2,6	0,0	(62,0)	3	199	\DRP14.\$SYSTEM.SYS00.NCPOBJ
3	\$YMIOP	0,5	1,5	(6,4)	80	205	\DRP14.\$SYSTEM.SYS00.TMIOP
5	\$Z0	0,7	1,7	(1,2)	102	200	\DRP14.\$SYSTEM.SYS00.OCDIST
6	\$SYSTEM	0,257	1,257	(3,45)	4096	220	\DRP14.\$SYSTEM.SYS00.TSYSDP2
7	\$ZOPR	0,8	1,8	(1,0)	102	201	\DRP14.\$SYSTEM.SYS00.OAUX
63	\$ZZKRN	0,294	1,328	(66,0)	4096	180	\DRP14.\$SYSTEM.SYS00.OZKRN
64	\$ZZWAN	0,291	1,298	(50,3)	132	180	\DRP14.\$SYSTEM.SYS00.WANMGR
65	\$ZZSTO	0,292	1,329	(65,0)	4096	180	\DRP14.\$SYSTEM.SYS00.TZSTO
66	\$ZZSMN	1,289	2,282	(64,1)	132	199	\DRP14.\$SYSTEM.SYS00.SANMAN
67	\$ZZSCL	1,290	2,277	(64,0)	132	199	\DRP14.\$SYSTEM.SYS00.SNETMON
68	\$ZZLAN	0,293	1,297	(43,0)	132	199	\DRP14.\$SYSTEM.SYS00.LANMAN
86	\$ZSNET	0,294	1,328	(66,0)	4096	180	\DRP14.\$SYSTEM.SYS00.OZKRN
87	\$ZSLM2	0,288	1,293	(67,0)	1024	221	\DRP14.\$SYSTEM.SYS00.TZSLM2
91	\$ZNET	0,14	1,13	(50,63)	3900	175	\DRP14.\$SYSTEM.SYS00.SCP
104	\$ZM03	3,279	0,0	(45,0)	132	201	\DRP14.\$SYSTEM.SYS00.QIOMON
105	\$ZM02	2,280	0,0	(45,0)	132	201	\DRP14.\$SYSTEM.SYS00.QIOMON
106	\$ZM01	1,280	0,0	(45,0)	132	201	\DRP14.\$SYSTEM.SYS00.QIOMON
107	\$ZM00	0,290	0,0	(45,0)	132	201	\DRP14.\$SYSTEM.SYS00.QIOMON
108	\$ZLOG	0,307	1,345	(1,0)	4024	150	\DRP14.\$SYSTEM.SYS00.EMSACOLL
104	\$ZM03	3,279	0,0	(45,0)	132	201	\DRP14.\$SYSTEM.SYS00.QIOMON
105	\$ZM02	2,280	0,0	(45,0)	132	201	\DRP14.\$SYSTEM.SYS00.QIOMON
106	\$ZM01	1,280	0,0	(45,0)	132	201	\DRP14.\$SYSTEM.SYS00.QIOMON
107	\$ZM00	0,290	0,0	(45,0)	132	201	\DRP14.\$SYSTEM.SYS00.QIOMON
108	\$ZLOG	0,307	1,345	(1,0)	4024	150	\DRP14.\$SYSTEM.SYS00.EMSACOLL
121	\$ZIM03	3,280	0,0	(64,2)	132	199	\DRP14.\$SYSTEM.SYS00.MSGMON
122	\$ZIM02	2,285	0,0	(64,2)	132	199	\DRP14.\$SYSTEM.SYS00.MSGMON
123	\$ZIM01	1,291	0,0	(64,2)	132	199	\DRP14.\$SYSTEM.SYS00.MSGMON
124	\$ZIM00	0,305	0,0	(64,2)	132	199	\DRP14.\$SYSTEM.SYS00.MSGMON
126	\$ZEXP	0,13	1,18	(63,30)	132	150	\DRP14.\$SYSTEM.SYS00.OZEXP
128	\$SC26	2,281	3,285	(63,4)	1	199	\DRP14.\$SYSTEM.SYS00.LHOBJ
129	\$SC25	2,283	3,286	(63,4)	1	199	\DRP14.\$SYSTEM.SYS00.LHOBJ
131	\$DATA6	0,296	1,287	(3,42)	4096	220	\DRP14.\$SYSTEM.SYS00.TSYSDP2
132	\$DATA5	0,297	1,286	(3,42)	4096	220	\DRP14.\$SYSTEM.SYS00.TSYSDP2
133	\$DATA4	0,298	1,285	(3,44)	4096	220	\DRP14.\$SYSTEM.SYS00.TSYSDP2
134	\$DATA3	0,299	1,284	(3,42)	4096	220	\DRP14.\$SYSTEM.SYS00.TSYSDP2
135	\$DATA2	0,300	1,283	(3,42)	4096	220	\DRP14.\$SYSTEM.SYS00.TSYSDP2
136	\$DATA1	0,301	1,282	(3,44)	4096	220	\DRP14.\$SYSTEM.SYS00.TSYSDP2
137	\$DATA	0,302	1,281	(3,44)	4096	220	\DRP14.\$SYSTEM.SYS00.TSYSDP2
145	\$ZOLHD	0,369	1,359	(1,30)	132	150	\DRP14.\$SYSTEM.SYS00.EMSDIST
167	\$ZTCO	0,338	1,332	(48,0)	32000	200	\DRP14.\$SYSTEM.SYS00.TCPIP
168	\$ZTNT	0,340	1,334	(46,0)	6144	149	\DRP14.\$SYSTEM.SYS00.TELSERV
200	\$ZPMON	0,375	0,0	(24,0)	4096	180	\DRP14.\$SYSTEM.SYS00.OSSMON

The columns in [Example 2-1](#) mean:

- LDev The logical device number
- Name The logical device name
- PPID The primary processor number and process identification number (PIN) of the specified device
- BPID The backup processor number and PIN of the specified device
- Type The device type and subtype
- RSize The record size the device is configured for
- Pri The priority level of the I/O process
- Program The fully qualified name of the program file for the process

[Table 2-1](#) gives the names of some subsystems that are common to most Integrity NonStop NS-series systems and are routinely monitored by operations. These subsystems appear in the LISTDEV output in [Example 2-1](#) on page 2-7.

Table 2-1. Key Subsystems and Their Logical Device Names and Device Types

Subsystem Name	Logical Name	Device Type	Description
TCP/IP	\$ZTCO	48	Transmission Control Protocol/Internet Protocol (TCP/IP)
Kernel	\$ZZKRN	66	NonStop Kernel operating system
Storage	\$ZZSTO	Disk: 3 Tape: 4 Open SCSI: 8 SMF pool: 25 SMF monitor: 52 \$ZZSTO: 65 \$ZSLM: 67	All storage devices; for example, disk and tape
SLSA	\$ZZLAN	43	All ServerNet LAN Systems Access (SLSA) connection and facilities
WAN	\$ZZWAN	50	All wide area network (WAN) connections

Also, in [Example 2-1](#) on [page 2-7](#), several disk drives and tape drives have been configured. You can identify the subsystem that owns a device by looking up its device type in the *SCF Reference Manual for H-Series RVUs*.

To display information about a particular device:

```
> SCF LISTDEV TYPE n
```

where *n* is a number for the device type. For example, if *n* is 3, the device type is disks. For the \MS9 system, entering LISTDEV TYPE 3 would display information for \$DATA6, \$DATA5, \$DATA4, \$DATA3, \$DATA2, \$DATA1, and \$DATA.

To display information for a given subsystem:

```
> SCF LISTDEV subsysname
```

where *subsysname* is the logical name of a subsystem; for example, \$ZZKRN for the Kernel subsystem.

Displaying SCF Configuration Information for Subsystems

The following tables give some of the SCF commands that display configuration information for objects controlled by subsystems that are common to most Integrity NonStop NS-series systems. The examples use the SCF ASSUME command to make a given subsystem the current default object for gathering information.

TCP/IP Subsystem

These examples are based on a TCP/IP process named \$ZTCO. Before using the commands listed in [Table 2-2](#), type this command to make the TCP/IP subsystem the default object:

```
> SCF ASSUME PROCESS $ZTCO
```

Table 2-2. Displaying Information for the TCP/IP Subsystem (\$ZTCO)

To Display Information About These Configured Objects	Enter This Command
All TCP/IP devices	LISTDEV TCPIP
Detailed information about the TCP/IP subsystem manager	INFO, DETAIL
All SUBNET names	INFO SUBNET *
All ROUTE names	INFO ROUTE *

Integrity NonStop servers support two versions of TCP/IP—NonStop TCP/IPv6 and NonStop TCP/IP. When you use the SCF LISTDEV and INFO commands, all current TCP/IP processes are displayed. For more information, refer to the *TCP/IPv6 Configuration and Management Manual* and the *TCP/IP Configuration and Management Manual*.

Kernel Subsystem

Before using commands listed in [Table 2-3](#), type this command to make the Kernel subsystem the default object:

```
> SCF ASSUME PROCESS $ZZKRN
```

Generic processes are part of the SCF Kernel subsystem. Generic processes can be created by the operating system or by a user. Examples of generic processes created by the operating system are the Kernel, SLSA, the storage subsystem, and WAN subsystem manager processes. Examples of generic processes created by a user are a Pathway program, a third-party program, or a user-written program that you configure to be controlled by the operating system. The \$ZPM persistence manager starts and monitors all generic processes.

Table 2-3. Displaying Information for the Kernel Subsystem (\$ZZKRN)

To Display Information About These Configured Objects	Enter This Command
The Kernel subsystem manager and ServerNet process names	LISTDEV KERNEL
All Kernel subsystem object and process names	NAMES \$ZZKRN
All generic processes	INFO *
Detailed information about a generic process	INFO # <i>generic-process</i> , DETAIL

Storage Subsystem

The storage subsystem manages disk and tape drives as well as SCSI and HP NonStop Storage Management Foundation (SMF) devices. Use the commands listed in [Table 2-4](#) to display desired information.

Table 2-4. Displaying Information for the Storage Subsystem (\$ZZSTO)

To Display Information About These Configured Objects	Enter This Command
All disk and tape drives (list)	LISTDEV STORAGE
All storage subsystem objects and processes (by name)	NAMES \$ZZSTO
All disk drives (list)	LISTDEV TYPE 3
All disk drives (summary information)	INFO DISK \$*
A specific disk drive (detailed information)	INFO DISK \$ <i>name</i> , DETAIL
All tape drives (list)	LISTDEV TYPE 4
All tape drives (summary information)	INFO TAPE \$*
A specific tape drive (detailed information)	INFO TAPE \$ <i>name</i> , DETAIL

When displaying configuration files for disk and tape devices in the storage subsystem, you can use the OBEYFORM option with the INFO command to display currently defined attribute values in the format that you would use to set up a configuration file. Each attribute appears as a syntactically correct configuration command.

For example, this command shows all the attributes for \$SYSTEM in OBEYFORM:

```
-> INFO DISK $SYSTEM, OBEYFORM
```

This output appears as shown in [Example 2-2](#).

Example 2-2. SCF ADD DISK Command Output

```
ADD DISK $SYSTEM , &
SENDTO STORAGE , &
  BACKUPCPU 1 , &
  HIGHPIN ON , &
  PRIMARYCPU 0 , &
  PROGRAM $SYSTEM.SYSTEM.TSYSDP2 , &
  STARTSTATE STARTED , &
  PRIMARYLOCATION (11,1,11) , &
  PRIMARYSAC IOMF.SAC-2.GRP-11.MOD-1.SLOT-50 , &
  MIRRORLOCATION (11,1,12) , &
  MIRRORSAC IOMF.SAC-1.GRP-11.MOD-1.SLOT-55 , &
  AUDITTRAILBUFFER 0 , &
  AUTOREVIVE OFF , &
  AUTOSTART ON , &
  CBPOOLLEN 1000 , &
  FSTCACHING OFF , &
  FULLCHECKPOINTS ENABLED , &
  HALTONERROR 1 , &
  LKIDLONGPOOLLEN 8 , &
  LKTABLESPACELEN 15 , &
  MAXLOCKSPEROCB 5000 , &
  MAXLOCKSPERTCB 5000 , &
  NONAUDITEDINSERT OFF , &
  NUMDISKPROCESSES 4 , &
  OSSCACHING ON , &
  PROTECTDIRECTORY SERIAL , &
  REVIVEBLOCKS 10 , &
  REVIVEINTERVAL 100 , &
  REVIVEPRIORITY 0 , &
  REVIVERATE 0 , &
  SERIALWRITES ENABLED
```

You can create a command file containing the output by using the OUT option of the INFO command. For details, see the *SCF Reference Manual for the Storage Subsystem*.

To get detailed configuration information in command format for all disks on the system, issue this command:

```
-> INFO DISK $* , OBEYFORM
```

To get detailed configuration information in command format for all tape drives on the system, issue this command:

```
-> INFO TAPE $*,OBEYFORM
```

ServerNet LAN Systems Access (SLSA) Subsystem

Before using commands listed in [Table 2-5](#), type this command to make the SLSA subsystem the default object:

```
> SCF ASSUME PROCESS $ZZLAN
```

The SLSA subsystem provides access to parallel LAN and WAN I/O for Integrity NonStop servers. The SLSA subsystem provides access to Ethernet, token-ring, and multifunction I/O board Ethernet adapters and to the ServerNet wide area network (SWAN) concentrator.

Table 2-5. Displaying Information for the SLSA Subsystem (\$ZZLAN)

To Display Information About These Configured Objects	Enter This Command
The SLSA subsystem manager	LISTDEV SLSA
All SLSA subsystem object and process names	NAMES \$ZZLAN
All configured adapters, with group/module/slot and adapter type	INFO ADAPTER *
A specific adapter	INFO ADAPTER <i>adapter</i> , DETAIL
All logical interface (LIF) names, with associated MAC addresses, associated physical interface (PIF) names, and port types	INFO LIF *
A specific LIF	INFO LIF <i>lifname</i> , DETAIL
A specific PIF	INFO PIF <i>pifname</i> , DETAIL
All ServerNet addressable controller (SAC) names	INFO SAC *
A specific SAC	INFO SAC <i>sacname.n</i> , DETAIL

When displaying configuration files for adapter and LIF devices in the SLSA subsystem, you can use the OBEYFORM option with the INFO command to display currently defined attribute values in the format that you would use to set up a configuration file. Each attribute appears as a syntactically correct system configuration command. For example:

```
ADD ADAPTER $ZZLAN.E0154, &
      LOCATION (1, 1, 54), &
      TYPE G4SA, &
      ACCESSLIST (0, 1)
```

Examples of the INFO command used with the OBEYFORM option are:

```
-> INFO ADAPTER $* , OBEYFORM
```

```
-> INFO LIF $* , OBEYFORM
```

WAN Subsystem

Before using commands listed in [Table 2-6](#), type this command to make the wide area network (WAN) subsystem the default object:

```
> SCF ASSUME PROCESS $ZZWAN
```

The WAN subsystem has responsibility for all WAN connections.

Table 2-6. Displaying Information for the WAN Subsystem (\$ZZWAN)

To Display Information About These Configured Objects	Enter This Command
The WAN subsystem manager	LISTDEV WAN
All WAN configuration managers, TCP/IP processes, and WANBoot processes	INFO *
All PATH names	INFO PATH *
The WAN adapters	INFO ADAPTER *
All DEVICE objects	INFO DEVICE *
All PROFILE objects	INFO PROFILE *

Additional Subsystems Controlled by SCF

[Table 2-7](#) lists the names associated with additional subsystems that can be controlled by SCF, along with its device types. You can use SCF commands to display the current attribute values for these objects.

Some SCF commands are available only to some subsystems. The objects that each command affects and the attributes of those objects are subsystem specific. This subsystem-specific information is presented in a separate manual for each subsystem. A partial list of these manuals appears in [Table 6-1](#) on page 6-13.

Refer to the *SCF Reference Manual for H-Series RVUs* for further information.

Table 2-7. Subsystem Objects Controlled by SCF (page 1 of 2)

Subsystem Acronym	Description	Device Type	Device Subtype
AM3270	AM3270 Access Method	60	0 or 10
ATM	Asynchronous Transfer Mode (ATM) protocol	42	0 or 1
ATP6100	Asynchronous Terminal Process 6100	53	0

Table 2-7. Subsystem Objects Controlled by SCF (page 2 of 2)

Subsystem Acronym	Description	Device Type	Device Subtype
CP6100	Communications Process Subsystem	51	0
Envoy	Byte-synchronous and asynchronous communications data link-level interface	7	0
EnvoyACP/XF	Byte-synchronous communications data link-level interface	11	40, 41, 42, 43
Expand	Expand network control process (\$NCP) or line-handler process	62 or 63	2, 3, 5, or 6
GDS	General Device Support	57	
OSIAPLMG	Open Systems Interconnection/Application Manager	55	20
OSIAS	Open Systems Interconnection/Application Services	55	1-5
OSICMIP	Open Systems Interconnection/Common Management Information Protocol	55	24
OSIFTAM	Open Systems Interconnection/File Transfer, Access, and Management	55	21 or 25
OSIMHS	Open Systems Interconnection/Message Handling System	55	11 or 12
OSITS	Open Systems Interconnection/Transport Services	55	55, 4
OSS	Open System Services	24	0
PAM	Port Access Method		
QIO	Queued I/O product	45	0
SCP	Subsystem Control Point	50	63
SCS	SQL Communications Subsystem	38	0
SNAX/APN	SNAX Advanced Peer Networking	58 or 13	0
SNAX/XF	SNAX Extended Facility	58 or 13	
SNAXAPC	SNAX Advanced Program Communication	13	10
SNAXCRE	SNAX Creator-2	18	0
SNAXHLS	SNAX High-Level Support	13	5
SNMP	Simple Network Management Protocol agent	31	0
TELSERV	TCP/IP TELNET product	46	0
TR3271	TR3271 Access Method	60	1 or 11
X25AM	X.25 Access Method	61	0

Displaying Configuration Information—SCF Examples

These examples show SCF commands that display subsystem configuration information, along with the information that is returned. These commands are not preceded by an ASSUME command.

To display all the processes running in the Kernel subsystem:

```
-> INFO PROC $ZZKRN.*
```

The system displays a listing similar to that shown in [Example 2-3](#):

Example 2-3. SCF INFO PROCESS Command Output

```
32-> INFO PROCESS $ZZKRN.*
```

```
NONSTOP KERNEL - Info PROCESS \DRP09.$ZZKRN
```

Symbolic Name	*Name	*Autorestart	*Program
CLCI-TACL	\$CLCI	10	\$SYSTEM.SYSTEM.TACL
OSM-APPSRVR	\$ZOSM	10	\$SYSTEM.SYSTEM.APPSRVR
OSM-CIMOM	\$ZCMOM	5	\$SYSTEM.SYSTEM.CIMOM
OSM-CONFLH-RD	\$ZOLHI	0	\$SYSTEM.SYSTEM.TACL
OSM-OEV	\$ZOEV	10	\$SYSTEM.SYSTEM.EVTMGR
QATRAK	\$TRAK	10	\$SYSTEM.SYSTOOLS.QATRAK
QIOMON	\$ZMnn	10	\$SYSTEM.SYSTEM.QIOMON
ROUT	\$ZLnn	10	\$SYSTEM.SYSTEM.ROUT
SCP	\$ZNET	10	\$SYSTEM.SYSTEM.SCP
SP-EVENT	\$ZSPE	5	\$SYSTEM.SYSTEM.ZSPE
TFDSHLP	\$ZTHnn	10	\$SYSTEM.SYSTEM.TFDSHLP
ZEXP	\$ZEXP	10	\$SYSTEM.SYSTEM.OZEXP
ZHOME	\$ZHOME	10	\$SYSTEM.SYSTEM.ZHOME
ZLOG	\$ZLOG	5	\$SYSTEM.SYSTEM.EMSACOLL
ZSLM2	\$ZSLM2	10	\$SYSTEM.SYSTEM.TZSLM2
ZZKRN	\$ZZKRN	10	\$SYSTEM.SYSTEM.OZKRN
ZZLAN	\$ZZLAN	10	\$SYSTEM.SYSTEM.LANMAN
ZZSTO	\$ZZSTO	10	\$SYSTEM.SYSTEM.TZSTO
ZZWAN	\$ZZWAN	10	\$SYSTEM.SYSTEM.WANMGR

To display a list of all SAC names with their associated owners and access lists:

```
-> info sac $zzlan.*
```

The system displays a listing similar to that shown in [Example 2-4](#):

Example 2-4. SCF INFO SAC Command Output

```
-> INFO SAC $ZZLAN.*
```

```
SLSA Info SAC
```

Name	Owner	*Access List
\$ZZLAN.E4SA0.0	3	(3,2,1,0)
\$ZZLAN.E4SA0.1	3	(3,2,1,0)
\$ZZLAN.E4SA52.0	0	(0,1)
\$ZZLAN.E4SA52.1	0	(0,1)
\$ZZLAN.FESA0.0	0	(0,1,2,3,4,5,6,7)

To display configuration attribute values for all the WAN subsystem configuration managers, TCP/IP processes, and WANBoot processes:

```
-> INFO PROCESS $ZZWAN.*
```

The system displays a listing similar to that shown in [Example 2-5](#):

Example 2-5. SCF INFO PROCESS \$ZZWAN Command Output

```
-> INFO PROCESS $ZZWAN.*
```

```
WAN MANAGER Detailed Info Process \DRP09.$ZZWAN.#ZTXAE
```

```
RecSize..... 0                      *Type..... ( 0,49)
Preferred Cpu..... 0                  Alternate Cpu..... 1
HOSTIP Address.... 172.031.145.090
*IOOBJECT..... $SYSTEM.SYS00.SNMPTMUX
TCPIP Name..... $ZTC02
```

```
WAN MANAGER Detailed Info Process \DRP09.$ZZWAN.#0
```

```
RecSize..... 0                      *Type..... (50,00)
Preferred Cpu..... 0                  Alternate Cpu..... N/A
*IOOBJECT..... $SYSTEM.SYS00.CONMGR
```

To display detailed information about an Expand line-handler process:

```
->INFO LINE $line-name, DETAIL
```

where *\$line-name* is the logical line-handler process name.

The system displays a listing similar to [Example 2-6](#) for Expand-Over-NAM and Expand-Over-ServerNet line-handler processes.

Example 2-6. SCF INFO LINE Command Output

```
-> INFO LINE $SC151, DETAIL
```

```
L2Protocol      Net^Nam  TimeFactor..... 1 *SpeedK..... NOT_SET
Framesize..... 132 -Rsize..... 1 -Speed.....
*LinePriority... 1 StartUp..... OFF Delay..... 0:00:00.10
*Rxwindow..... 7 *Timerbind... 0:01:00.00 *L2Timeout.... 0:00:01.00
*Txwindow..... 7 *Maxreconnects... 0 *AfterMaxRetries PASSIVE
*Timerreconnect 0:01:00.00 *Retryprobe..... 10 *Timerprobe... 0:00:30.00
*Associateddev... $ZZSCL *Associatesubdev *Timerinactivity 0:00:00.00
*ConnectType.... ACTIVEANDPASSIVE
*LineTf..... 0
```

3

Overview of Monitoring and Recovery

[When to Use This Section](#) on page 3-1

[Functions of Monitoring](#) on page 3-2

[Monitoring Tasks](#) on page 3-2

[Working With a Daily Checklist](#) on page 3-2

[Tools for Checking the Status of System Hardware](#) on page 3-3

[Additional Monitoring Tasks](#) on page 3-6

[Monitoring and Resolving Problems—An Approach](#) on page 3-7

[Using OSM to Monitor the System](#) on page 3-7

[Using the OSM Service Connection](#) on page 3-7

[Recovery Operations for Problems Detected by OSM](#) on page 3-12

[Monitoring Problem Incident Reports](#) on page 3-12

[Using SCF to Monitor the System](#) on page 3-12

[Determining Device States](#) on page 3-13

[Automating Routine System Monitoring](#) on page 3-16

[Using the Status LEDs to Monitor the System](#) on page 3-20

[Related Reading](#) on page 3-22

When to Use This Section

This section provides an overview of monitoring an Integrity NonStop server using various tools. It describes some common monitoring tasks. It also refers you to other sections or manuals for more information about monitoring specific system components, events, applications, or processes.

Functions of Monitoring

You must monitor a system to ensure that it is operating properly and to recognize when corrective action is required. By monitoring a system, you can:

- Verify whether components are currently up or down
- Be quickly notified of error conditions, state changes, and threshold conditions that have been exceeded or are reaching their limits
- View a chronological list of events that can help with problem diagnosis and resolution
- Determine how much of a particular resource is being used; for example, processor capacity, disk or file space, or communications line bandwidth
- Find performance problems that can affect the users of the system
- Make better use of existing resources
- Ensure that products such as HP NonStop SQL/MP, HP NonStop SQL/MX, HP NonStop Transaction Management Facility (TMF), and Pathway are available
- Prevent many problems and outages from occurring

Monitoring Tasks

Regardless of the shift you work, certain areas of your hardware and software environment need to be checked on a regular basis. This subsection provides guidelines that will enable you to determine the general areas you should monitor.

Working With a Daily Checklist

A good method for ensuring that certain areas of your operations environment are monitored is to develop a checklist. Monitor these items on a system frequently. At least daily, monitor:

- OSM Service Connection GUI
- Event messages
- Alarms
- Problem incident reports
- The status of all system components
- The status of processes
- The status of all applications
- The performance of processors, disks, and communications lines (Monitoring performance is not discussed in this guide.)

An example of a checklist you might use to standardize your routine daily monitoring tasks is:

Task	Operator's name	Date & time	Notes and questions
Check phone messages			
Check faxes			
Check e-mail			
Check shift log			
Check EMS event messages			
Check status of terminals			
Check comm. lines			
Check TMF status			
Check Pathway status			
Check disks			
Check tape drives			
Check processors			
Check printers			
Check spooler supervisor and collector processes			
Check ServerNet cluster status			

Tools for Checking the Status of System Hardware

Several tools are available to check the status of system components in an Integrity NonStop NS-series server. The most frequently used tools are the OSM Service Connection and the Subsystem Control Facility (SCF).

For information relating to system components in NonStop S-series servers, refer to the appropriate NonStop S-Series documentation.

[Table 3-1](#) lists the tools available to monitor system components.

Table 3-1. Monitoring System Components (page 1 of 3)

Resource	Monitored Using These Tools	See...
Adapters for communications subsystems: G4SA	OSM Service Connection SCF interface to various subsystems	Using the OSM Service Connection on page 3-7 Section 6, Communications Subsystems: Monitoring and Recovery Section 8, I/O Adapters and Modules: Monitoring and Recovery <i>OSM Service Connection User's Guide</i> (or OSM Service Connection online help)
Adapters for the storage subsystem: Fibre Channel ServerNet adapter (FCSA)	OSM Service Connection SCF interface to the storage subsystem	Using the OSM Service Connection on page 3-7 Section 8, I/O Adapters and Modules: Monitoring and Recovery <i>OSM Service Connection User's Guide</i> (or OSM Service Connection online help)
AWAN access server	RAS management tool	<i>AWAN 3886 Server Installation and Configuration Guide</i>
Communications lines	SCF interface to the various subsystems	Section 6, Communications Subsystems: Monitoring and Recovery
Disk drive enclosure (and individual disk drives) attached to FCSAs	OSM Service Connection SCF interface to the storage subsystem DSAP	Using the OSM Service Connection on page 3-7 Section 8, I/O Adapters and Modules: Monitoring and Recovery Section 10, Disk Drives: Monitoring and Recovery <i>Guardian User's Guide.</i>
Disk drives attached to ServerNet adapters in legacy NonStop S-series enclosures	OSM Service Connection SCF interface to the storage subsystem DSAP	Using the OSM Service Connection on page 3-7 Section 10, Disk Drives: Monitoring and Recovery <i>Guardian User's Guide</i>
Modular I/O adapter module (IOAM) and subcomponents, including ServerNet switch boards, power supplies, and fans	OSM Service Connection	Using the OSM Service Connection on page 3-7 Monitor Batteries on page 14-4 <i>OSM Service Connection User's Guide</i> (or OSM Service Connection online help)

Table 3-1. Monitoring System Components (page 2 of 3)

Resource	Monitored Using These Tools	See...
Legacy NonStop S-series enclosure and subcomponents, including IOMF2 CRUs, PMCUs, power supplies, fans, and batteries	OSM Service Connection	Using the OSM Service Connection on page 3-7 Section 8, I/O Adapters and Modules: Monitoring and Recovery <i>OSM Service Connection User's Guide</i> (or OSM Service Connection online help)
NonStop Blade Complex components: Blade Elements, LSUs, logical processors	OSM Service Connection	Using the OSM Service Connection on page 3-7 Section 9, Processors and Components: Monitoring and Recovery <i>OSM Service Connection User's Guide</i> (or OSM Service Connection online help)
NonStop ServerNet Cluster 6770 Switch	OSM Service Connection	<i>ServerNet Cluster 6770 Hardware Installation and Support Guide</i> , or <i>ServerNet Cluster Manual</i> <i>OSM Service Connection User's Guide</i> (or OSM Service Connection online help)
NonStop ServerNet Cluster 6780 Switch	OSM Service Connection	<i>ServerNet Cluster 6780 Operations Guide</i> <i>OSM Service Connection User's Guide</i> (or OSM Service Connection online help)
Printers	SCF SPOOLCOM	Section 12, Printers and Terminals: Monitoring and Recovery <i>Guardian User's Guide</i>
Processor switch (P-switch) module and subcomponents, including ServerNet switch boards, power supplies, fans, PICs and ports	OSM Service Connection	Using the OSM Service Connection on page 3-7 <i>OSM Service Connection User's Guide</i> (or OSM Service Connection online help)
ServerNet connectivity for an Integrity NonStop NS14000 or NS1000 system (which have no processor switches) 4-Port ServerNet Extender (4PSE)	OSM Service Connection	Using the OSM Service Connection on page 3-7 <i>OSM Service Connection User's Guide</i> (or OSM Service Connection online help)
ServerNet fabrics: processor-to-processor and processor-to-IOMF2 communication	OSM Service Connection SCF interface to the Kernel subsystem	Using the OSM Service Connection on page 3-7 Section 7, ServerNet Resources: Monitoring and Recovery

Table 3-1. Monitoring System Components (page 3 of 3)

Resource	Monitored Using These Tools	See...
ServerNet wide area network (SWAN) concentrator	OSM Service Connection SCF interface to the WAN subsystem	Using the OSM Service Connection on page 3-7 Section 6, Communications Subsystems: Monitoring and Recovery
Tape drives	OSM Service Connection SCF interface to the storage subsystem MEDIACOM	Section 11, Tape Drives: Monitoring and Recovery Section 8, I/O Adapters and Modules: Monitoring and Recovery <i>Guardian User's Guide</i>
Uninterruptible Power Supply (UPS)	OSM Service Connection	Monitor Batteries on page 14-4 <i>OSM Service Connection User's Guide</i> (or OSM Service Connection online help)

Additional Monitoring Tasks

[Table 3-2](#) provides an example of additional areas you should monitor daily.

Table 3-2. Daily Tasks Checklist

General Tasks	Specific Tasks	For More Information, See
Monitor messages from system users	Check telephone, fax, electronic mail, and any other messages	<i>Guardian User's Guide</i>
Monitor operator messages	From the OSM Event Viewer From the EMSDIST printing distributor From ViewPoint	Section 4, Monitoring EMS Event Messages OSM Event Viewer online help <i>Guardian User's Guide</i> <i>ViewPoint Manual</i>
Monitor key applications	Monitor Pathway and TMF Monitor SQL/MX, SQL/MP and other applications	Section 13, Applications: Monitoring and Recovery The documentation specific to the application
Monitor system processes	Use the SCF and TACL PPD commands	Section 5, Processes: Monitoring and Recovery

Monitoring and Resolving Problems—An Approach

A useful approach to identifying and resolving problems in your system is to first use OSM to locate the focal point of a hardware problem and then use SCF to gather all the related data from the subsystems that control or act on the hardware. In this way, you can develop a larger picture that encompasses the whole environment, including communications links and other objects and services that might be contributing to the problem or affected by it.

To get comprehensive online descriptions of all the available SCF commands, use the SCF HELP command.

The following subsections give instructions for using OSM and SCF to monitor and resolve problems.

Using OSM to Monitor the System

This section deals mostly with the OSM Service Connection, the primary OSM interface for system monitoring and serviceability.

See [Overview of OSM Applications](#) on page 1-11 for examples of how the other OSM applications are used for monitoring-related functions.

Using the OSM Service Connection

The OSM Service Connection can be used in a variety of ways to monitor your system, including:

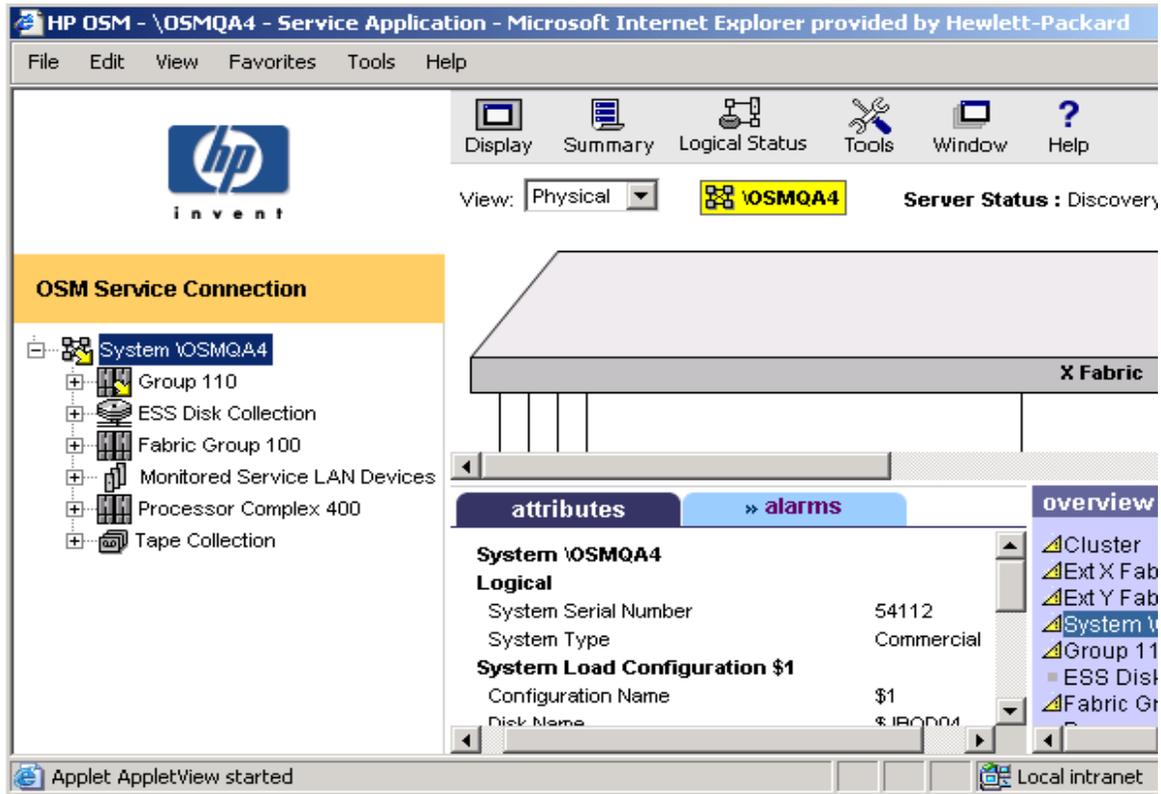
- Use of colors and symbols to direct you to the source of any problems
- Attribute values for system resources, displayed in the Attributes tab and in many dialog boxes.
- Alarms, displayed in the Alarms tab and Alarm Summary dialog box.

The following section presents one model for using the OSM Service Connection to monitor your system, along with a few other options.

A Top-Down Approach

The Management (or main) window of the OSM Service Connection uses a series of colors and symbols to notify you that problems exist within the system. You can tell at a high-level glance when problem conditions exist, then drill-down, or expand the tree pane to find the component reporting the problem. [Figure 3-1](#) illustrates how both the the rectangular system icon (located at the top of the view pane) and the system object in the tree pane indicate problems within the system. The system icon, which is green when OSM is reporting no problems on the system, has turned yellow. The system icon in the tree pane is displaying a yellow arrow to indicate a problem within.

Figure 3-1. OSM Management: System Icons Indicate Problems Within



VST310.vsd

Note. In the OSM Service Connection Management window, the tree pane is located on the far left. In the lower right is the Overview pane. Located between them is the details pane, from which you can choose to view the Attributes or Alarms tab. Directly above the details pane is the view pane, from which you can choose a Physical or Inventory view of your system or ServerNet Cluster. The gray bar directly above the view pane is an OSM-specific toolbar (as opposed to the standard Internet Explorer menu bar at the top of the browser window).

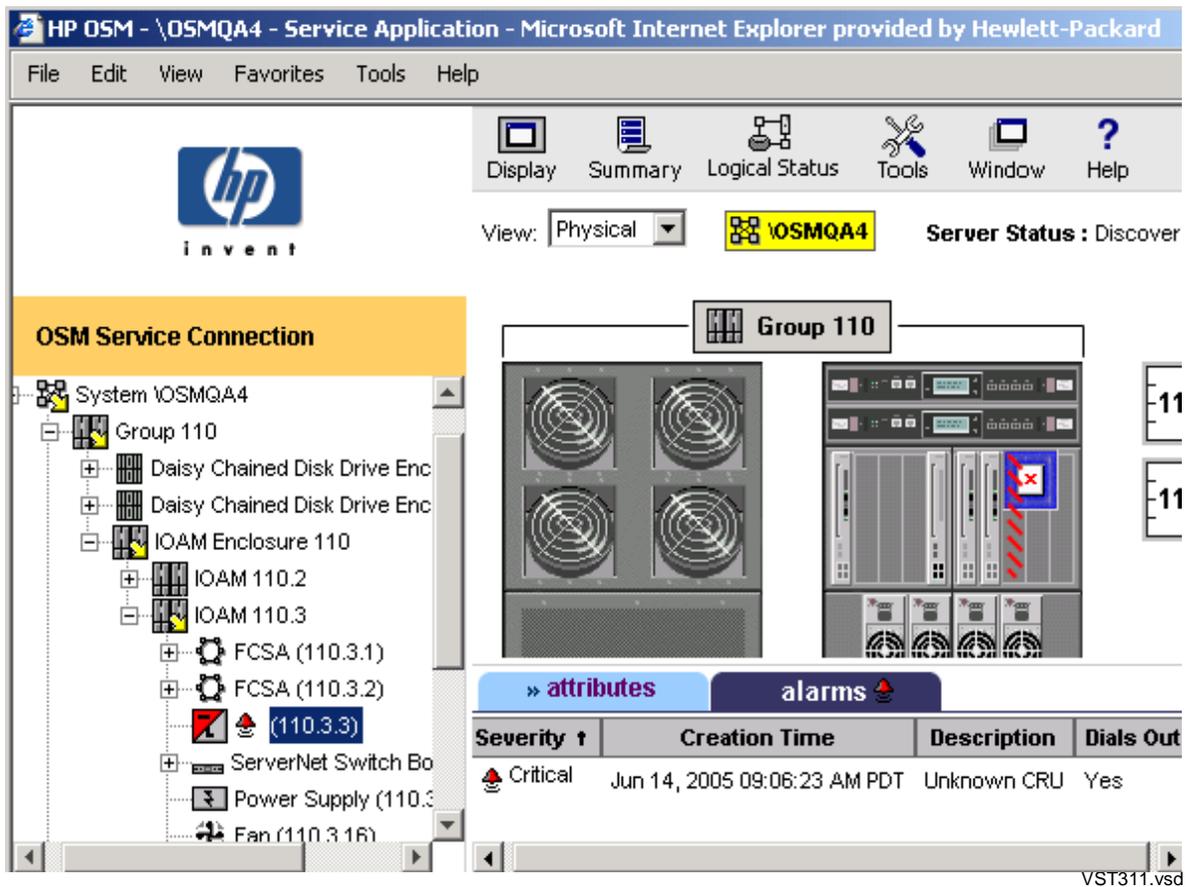
Expanding the system object in the tree pane, you can see a yellow arrow on the Group 110 object, indicating that the problem is located somewhere within that group.

Expanding the tree pane further, as illustrated in [Figure 3-2](#), yellow arrows on the IOAM Enclosure 110 and IOAM 110.3 objects reveal that the problem exists on a ServerNet adapter in slot 3 of that I/O module. The red bell-shaped icon by that resource object (in the tree pane) indicates that there is an alarm on the object. To obtain information about the alarm:

1. Click to select the object displaying the red triangular and bell-shaped symbols.
2. Select the Alarms tab from the details pane.

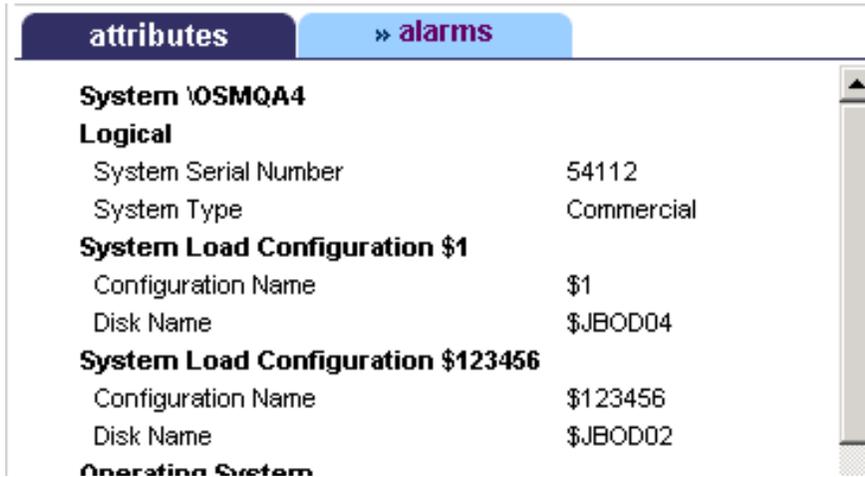
3. Click to select the alarm, then right-click and select Details.

Figure 3-2. Expanding the Tree Pane to Locate the Source of Problems



Check the Attributes tab (Figure 3-3) also, as a yellow or red triangular symbol indicates problem attribute values exist. In this case, the degraded Service State attribute was caused by an alarm. However, when a resource displays a yellow or red triangular object but no bell-shaped icon, it has no alarms but is reporting problem or degraded attribute values.

Figure 3-3. Attributes Tab



VST312.vsd

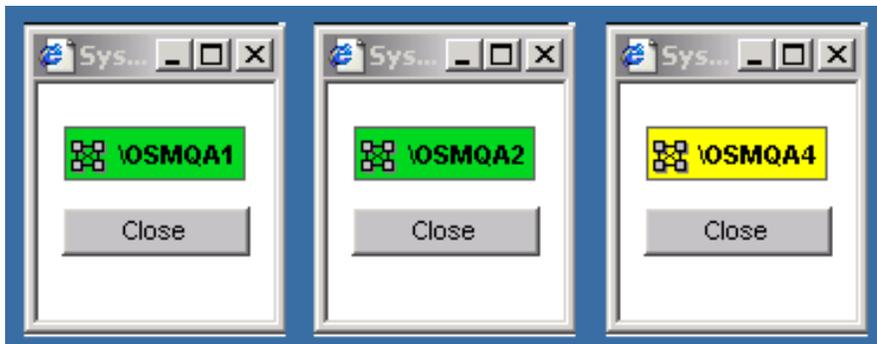
Using System Status Icons to Monitor Multiple Systems

When you are monitoring multiple systems, you can create a System Status Icon for each system, allowing you to keep a high-level eye on each system while saving screen space. [Figure 3-4](#) shows three separate System Status icons, each created by:

1. Establishing an OSM Service Connection session to the system.
2. From the Summary menu on the OSM toolbar, selecting System Status.

You can then minimize, but not close, the OSM Service Connection Management window for each system. If the System Status icon for a system turns from green to yellow, as illustrated in [Figure 3-4](#), open the Management window for that system and locate the problem as described in [A Top-Down Approach](#) on page 3-7.

Figure 3-4. Using System Status Icons to Monitor Multiple Systems



VST313.vsd

Using Alarm and Problem Summaries

Other options for monitoring your system with the OSM Service Connection include using the Alarm Summary (Figure 3-5) or Problem Summary (Figure 3-6) dialog boxes to quickly view all alarms and problem conditions that exist on your system.

Figure 3-5. Alarm Summary Dialog Box

Severity ↑	Creation Time	IndicationIdentifier	Description	Probable Cause
Critical	Jun 16, 2005 09:10:19 PM	(100.2.2)	Unknown CRU	Equipment Malfunction
Critical	Jun 16, 2005 09:10:19 PM	Port (100.2.5.3)	Neighbor Check Timed Out For Port	Connection Establishment Error
Critical	Jun 16, 2005 09:10:19 PM	Port (100.2.5.4)	Neighbor Check Timed Out For Port	Connection Establishment Error
Critical	Jun 16, 2005 09:10:19 PM	(100.3.2)	Unknown CRU	Equipment Malfunction
Critical	Jun 16, 2005 09:10:19 PM	Port (100.3.5.4)	Neighbor Check Timed Out For Port	Connection Establishment Error
Major	Jun 16, 2005 07:56:23 PM	Storage Router \$EXT9	Not Communicating	Equipment Malfunction
Major	Jun 16, 2005 07:56:23 PM	Storage Router \$EXT8	Not Communicating	Equipment Malfunction
Major	Jun 16, 2005 09:10:19 PM	Port (100.2.4.1)	ServerNet Port is Disconnected	Connection Establishment Error
Major	Jun 16, 2005 09:10:19 PM	Processor Switch Power Supply	AC Power Loss	Power Problem
Major	Jun 16, 2005 09:10:19 PM	Processor Switch Power Supply	AC Power Loss	Power Problem
Minor	Jun 16, 2005 07:56:23 PM	Tape Drive \$TAPE1	Storage Router Not Supported	Configuration/Customization Error

Detail... Delete Suppress Unsuppress Close Help

VST314.vsd

Figure 3-6. Problem Summary Dialog Box

Name	Service State ↓	Problem Attribute Category	Attributes In Problem
Power Supply (110.211.1.95)	Service Required	Logical Physical	AC Power State
FC Disk \$JBOD02-P (110.211.1.3)	Service Required	Physical	
(110.3.3)	Service Required	Physical	
UPS rx5500	Problem Acknowledged	Physical	
Processor Switch Power Supply (100.3.18)	Problem Acknowledged	Logical	
Processor Switch Power Supply (100.2.18)	Problem Acknowledged	Logical	
Port (100.3.10.4)	Problem Acknowledged	Logical	
Port (100.2.10.4)	Problem Acknowledged	Logical	
Tape Drive \$VTAPE06	OK	Logical	
Processor Switch Power Supply (100.2.18)	Problem Acknowledged	Logical	
Port (100.3.10.4)	Problem Acknowledged	Logical	

Close Help

VST315.vsd

Suppressing Problems and Alarms

In certain cases, you might want to acknowledge or suppress a particular problem, to stop it from propagating a known problem all the way up to the system level. That way, it will be easier to identify other problems that might occur. For more information on OSM problem management features such as deleting or suppressing alarms and suppressing problem attributes, see the *OSM Service Connection User's Guide* (also available as online help within the OSM Service Connection).

Recovery Operations for Problems Detected by OSM

Recovery operations depend on the particular problem, of course. Methods of determining the appropriate recovery action include:

- Alarm Details, available for each alarm displayed in OSM, provide suggested repair actions.
- The value displayed by problem attributes in OSM often provide clues to recovery.
- EMS events, retrieved and viewed in the OSM Event Viewer, include cause, effect, and recovery information in the event details.
- Check the section in this guide that covers the system resource—for example, [Section 11, Tape Drives: Monitoring and Recovery](#)—for information on using the SCF and other tools to determine the cause of a problem. Then follow the directions in the Recovery Operations subsection in the relevant section.

Replacing a system component that has malfunctioned is beyond the scope of this guide. For more information, contact your service provider, or refer to the [Support and Service Library](#) on page 1-12.

Monitoring Problem Incident Reports

The OSM Notification Director generates problem incident reports when changes occur that could directly affect the availability of resources on your Integrity NonStop server. The Incident Report List tab on the Notification Director dialog box allows you to view, sort, authorize, and reject incident reports. The Notification Director allows you to forward notifications to your service provider if your system is configured for remote dial-out.

Using SCF to Monitor the System

Use the Subsystem Control Facility (SCF) to display information and current status for all the devices on your system known to SCF. Some SCF commands are available only to some subsystems. The objects that each command affects and the attributes of those objects are subsystem specific. This subsystem-specific information appears in a separate manual for each subsystem. A partial list of these manuals appears in [Appendix C, Related Reading](#).

Determining Device States

This subsection explains how to determine the state of devices on your system. For example, to monitor the current state of all tape devices on your system, at an SCF prompt:

```
-> STATUS TAPE $*
```

[Example 3-1](#) shows the results of the SCF STATUS TAPE \$* command:

Example 3-1. SCF STATUS TAPE Command

```
1-> STATUS TAPE $*
STORAGE - Status TAPE \COMM.$TAPE0
LDev  State      Primary  Backup  DeviceStatus
      State      PID      PID
 156   STOPPED    2,268   3,288   NOT READY

STORAGE - Status TAPE \COMM.$DLT20
LDev  State      Primary  Backup  DeviceStatus
      State      PID      PID
 394   STARTED    2,267   3,295   NOT READY

STORAGE - Status TAPE \COMM.$DLT21
LDev  State      Primary  Backup  DeviceStatus
      State      PID      PID
 393   STARTED    1,289   0,299   NOT READY

STORAGE - Status TAPE \COMM.$DLT22
LDev  State      Primary  Backup  DeviceStatus
      State      PID      PID
 392   STARTED    0,300   1,288   NOT READY

STORAGE - Status TAPE \COMM.$DLT23
LDev  State      Primary  Backup  DeviceStatus
      State      PID      PID
 391   STARTED    1,287   0,301   NOT READY

STORAGE - Status TAPE \COMM.$DLT24
LDev  State      Primary  Backup  DeviceStatus
      State      PID      PID
 390   STARTED    6,265   7,298   NOT READY

STORAGE - Status TAPE \COMM.$DLT25
LDev  State      Primary  Backup  DeviceStatus
      State      PID      PID
 389   STARTED    4,265   5,285   NOT READY
```

Some other examples of the SCF STATUS command are:

```
-> STATUS LINE $LAM3
-> STATUS WS $LAM3.#WS1
-> STATUS WS $LAM3.*
-> STATUS WINDOW $LAM3.#WS1.*
-> STATUS WINDOW $LAM3.*, SEL STOPPED
```

The general format of the STATUS display follows. However, the format varies depending on the subsystem.

<i>subsystem STATUS object-type object-name</i>							
Name	State	PPID	BPID	<i>attr1</i>	<i>attr2</i>	<i>attr3</i>	...
object-name1	state	nn,nnn	nn,nnn	val1	val2	val3	...
object-name2	state	nn,nnn	nn,nnn	val1	val2	val3	...

where:

<i>subsystem</i>	The reporting subsystem name
<i>object-type</i>	The object, or device, type
<i>object-name</i>	The fully qualified name of the object
State	One of the valid object states: ABORTING, DEFINED, DIAGNOSING, INITIALIZED, SERVICING, STARTED, STARTING, STOPPED, STOPPING, SUSPENDED, SUSPENDING, and UNKNOWN
PPID	The primary processor number and process identification number (PIN) of the object
BPID	The backup processor number and PIN of the object
<i>attrn</i>	The name of an attribute of the object
<i>valn</i>	The value of that object attribute

SCF Object States

[Table 3-3](#) lists and explains the possible object states that the SCF STATUS command can report.

Table 3-3. SCF Object States (page 1 of 2)

State	Substate	Explanation
ABORTING		The object is being aborted. The object is responding to an ABORT command or some type of malfunction. In this state, no new links are allowed, and drastic measures might be underway to reach the STOPPED state. This state is irrevocable.
DEFINED		One of the generally defined possible conditions of an object with respect to the management of that object.
DIAGNOSING		The object is in a subsystem-defined test mode entered through the DIAGNOSE command.
INITIALIZED		The system has created the process, but it is not yet in one of the operational states.

Table 3-3. SCF Object States (page 2 of 2)

State	Substate	Explanation
SERVICING	SPECIAL	The object is being serviced or used by a privileged process and is inaccessible to user processes.
	TEST	The object is reserved for exclusive testing.
STARTED		The object is logically accessible to user processes.
STARTING		The object is being initialized and is in transition to the STARTED state.
STOPPED	CONFIG-ERROR	The object is configured improperly.
	DOWN	The object is no longer logically accessible to user processes.
	HARDDOWN	The object is in the hard-down state or is physically inaccessible due to a hardware error.
	INACCESSIBLE	The object is inaccessible to user processes.
	PREMATURE-TAKEOVER	The backup input/output (I/O) process was asked to take over for the primary I/O process before it had the proper information.
	RESOURCE-UNAVAILABLE	The input/output (I/O) process could not obtain a necessary resource.
STOPPING	UNKNOWN-REASON	The input/output (I/O) process is down for an unknown reason.
		The object is in transition to the STOPPED state. No new links are allowed to or from the object. Existing links are in the process of being deleted.
SUSPENDED		The flow of information to and from the object is restricted. (It is typically prevented.) A subsystem must clearly distinguish between the type of information that is allowed to flow in the SUSPENDED state and that which normally flows in the STARTED or STOPPED state. In the SUSPENDED state, the object must complete any outstanding work defined by the subsystem.
SUSPENDING		The object is in transition to the SUSPENDED state. The subsystem must clearly define the nature of the restrictions that this state imposes on its objects.
UNKNOWN		The object's state cannot be determined because the object is inaccessible.

Automating Routine System Monitoring

You can automate many of the monitoring procedures. Automation saves you time and helps you to perform many routine tasks more efficiently.

Your operations environment might be using TACL macros, TACL routines, or command files to perform routine system monitoring and other tasks. These items allow you to run many procedures so that you can quickly determine system status, produce reports, or perform other common tasks. The *TACL Reference Manual* contains an example that you can adapt to automate system monitoring.

[Example 3-2](#) contains an example of a command file you can use or adapt to check many of the system elements:

1. To create a command file named SYSCHK that will automate system monitoring, type the text shown in [Example 3-2](#) into an EDIT file.

Example 3-2. System Monitoring Command File

```
COMMENT THIS IS THE FILE SYSCHK
COMMENT THIS CHECKS ALL DISKS:
SCF STATUS DISK $*
COMMENT THIS CHECKS ALL TAPE DRIVES:
SCF STATUS TAPE $*
COMMENT THIS CHECKS THE SPOOLER PRINT DEVICES:
SPOOLCOM DEV
COMMENT THIS CHECKS THE LINE HANDLERS:
SCF STATUS LINE $*
COMMENT THIS CHECKS THE STATUS OF TMF:
TMFCOM;STATUS TMF
COMMENT THIS CHECKS THE STATUS OF PATHWAY:
PATHCOM $ZVPT;STATUS PATHWAY;STATUS PATHMON
COMMENT THIS CHECKS ALL SACS:
SCF STATUS SAC $*
COMMENT THIS CHECKS ALL ADAPTERS
SCF STATUS ADAPTER $*
COMMENT THIS CHECKS ALL LIFS
SCF STATUS LIF $*
COMMENT THIS CHECKS ALL PIFS
SCF STATUS PIF $*
```

2. After you create this file, at a TACL prompt, type this command to execute the file and automatically monitor many elements of your system:

```
> OBEY SYSCHK
```

For an example of the output that is sent to your home terminal when you execute a command file such as SYSCHK, refer to [Example 3-3](#). This output shows that all elements of the system being monitored are up and running normally.

Example 3-3. System Monitoring Output File (page 1 of 3)

```

COMMENT THIS IS THE FILE SYSCHK

COMMENT THIS CHECKS ALL DISKS:
SCF STATUS DISK $*

STORAGE - Status DISK \SHARK.$DATA12
LDev Primary Backup Mirror MirrorBackup Primary Backup
          PID          PID
    52 *STARTED   STARTED *STARTED   STARTED   3,262   2,263

STORAGE - Status DISK \SHARK.$DATA01
LDev Primary Backup Mirror MirrorBackup Primary Backup
          PID          PID
    63 *STARTED   STARTED *STARTED   STARTED   0,267   1,266

STORAGE - Status DISK \SHARK.$DATA04
LDev Primary Backup Mirror MirrorBackup Primary Backup
          PID          PID
    60 *STARTED   STARTED *STARTED   STARTED   0,270   1,263

STORAGE - Status DISK \SHARK.$SYSTEM
LDev Primary Backup Mirror MirrorBackup Primary Backup
          PID          PID
    6  *STARTED   STARTED  STOPPED   STOPPED   0,256   1,256

COMMENT THIS CHECKS ALL TAPE DRIVES:
SCF STATUS TAPE $*

STORAGE - Status TAPE $TAPE1
LDev State SubState Primary Backup DeviceStatus
          PID          PID
    48  STARTED          0,274

STORAGE - Status TAPE $TAPE0
LDev State SubState Primary Backup DeviceStatus
          PID          PID
    49  STARTED          0,273

COMMENT THIS CHECKS THE SPOOLER PRINT DEVICES:
SPOOLCOM DEV

DEVICE STATE FLAGS PROC FORM
$LINE1 WAITING H $SPLX
$LINE2 WAITING H $SPLX
$LINE3 WAITING H $SPLX
$LASER WAITING H $SPLP
    
```

Example 3-3. System Monitoring Output File (page 2 of 3)

COMMENT THIS CHECKS ALL SACS:
SCF STATUS SAC \$*

SLSA Status SAC

Name	Owner	State
\$ZZLAN.E4SA1.0	1	STARTED
\$ZZLAN.E4SA1.1	0	STARTED
\$ZZLAN.E4SA1.2	0	STARTED
\$ZZLAN.E4SA1.3	1	STARTED

COMMENT THIS CHECKS ALL ADAPTERS
SCF STATUS ADAPTER \$*

SLSA Status ADAPTER

Name	State
\$ZZLAN.MIOE0	STARTED
\$ZZLAN.E4SA0	STARTED
\$ZZLAN.MIOE1	STARTED
\$ZZLAN.E4SA2	STARTED

COMMENT THIS CHECKS ALL LIFS
SCF STATUS LIF \$*

SLSA Status LIF

Name	State	Access State
\$ZZLAN.LAN0	STARTED	UP
\$ZZLAN.LAN3	STARTED	DOWN

COMMENT THIS CHECKS ALL PIFS
SCF STATUS PIF \$*

SLSA Status PIF

Name	State
\$ZZLAN.E4SA0.0.A	STARTED
\$ZZLAN.E4SA0.0.B	STARTED
\$ZZLAN.E4SA0.1.A	STOPPED
\$ZZLAN.E4SA0.1.B	STARTED

COMMENT THIS CHECKS THE LINE HANDLERS:
SCF STATUS LINE \$*

COMMENT THIS CHECKS THE STATUS OF TMF:
TMFCOM:STATUS TMF

TMF Status:

System: \SAGE, Time: 12-Jul-1994 14:05:00

State: Started

Transaction Rate: 0.25 TPS

AuditTrail Status:

Master:

Active audit trail capacity used: 68%

First pinned file: \$TMF1.ZTMFAT.AA000044

Reason: Active transactions(s).

Current file: \$TMF1.ZTMFAT.AA000045

AuditDump Status:

Master: State: enabled, Status: active, Process \$X545,

File: \$TMF2.ZTMFAT.AA000042

BeginTrans Status: Enabled

Catalog Status:

Status: Up

Processes Status:

Dump Files:

#0: State: InProgress

Example 3-3. System Monitoring Output File (page 3 of 3)

```

COMMENT THIS CHECKS THE STATUS OF PATHWAY:
PATHCOM $ZVPT;STATUS PATHWAY;STATUS PATHMON

PATHWAY -- STATE=RUNNING

                RUNNING
EXTERNALTCPS      0
LINKMONS          0
PATHCOMS          1
SPI               0

                                FREEZE
                RUNNING  STOPPED  THAWED   FROZEN  PENDING
SERVERCLASSES    17      0        17      0       0

                RUNNING  STOPPED  PENDING
SERVERPROCESSES  17      35      0
TCPS             1       0       0

                RUNNING  STOPPED  PENDING  SUSPENDED
TERMS            0       0       0        0
PATHMON \COMM.$ZVPT -- STATE=RUNNING          CPUS 0:1
  PATHCTL (OPEN)  $OPER.VIEWPT.PATHCTL
  LOG1 SE (OPEN)  $0
  LOG2 (CLOSED)

REQNUM  FILE      PID          PAID    WAIT
1       PATHCOM  $X0X7      1,254
2       TCP      $Z040
    
```

Using the Status LEDs to Monitor the System

Status LEDs on the various enclosures and system components light during certain operations, such as when the system performs a series of power-on self-tests (POSTs) when a server is first powered on. [Table 3-4](#) lists some of the status light-emitting diodes (LEDs) and their functions.

Table 3-4. Status LEDs and Their Functions (page 1 of 3)

Location	LED Name	Color	Function
Disk drive	Power-on	Green	Lights when the disk drive is receiving power.
	Activity	Yellow or amber	Lights when the disk drive is executing a read or write command.
Disk drive, fibre channel	Drive Ready (top green)	Green	Flashes when drive is starting. (At the same time, the middle green light is lit and the bottom amber light is lit.)
	Drive Online (middle green)	Green	Flashes when drive is operational and performing a locate function.
	Drive Failure (bottom amber)	Amber	Flashes when drive is inactive or in error condition. When this occurs, verify the loop and replace the drive, if necessary.
	All		If all lights are on and none are flashing, the drive is not operational. Perform the following actions: <ol style="list-style-type: none"> 1. Check FC-SA. Replace if defective. 2. Check FC-AL I/O module. Replace if defective. 3. Replace drive.
EMU	Heartbeat	Left Green	Flashes when EMU is operational and performing locate. Power might just have been applied to the EMU, or an enclosure fault might exist. On when an EMU fault exists that is not an enclosure fault. Off when an EMU fault exists, which could be or might not be an enclosure fault.
	Power	Middle Green	Flashes when EMU is operational and performing locate. On when EMU is operational. An EMU or an enclosure fault might still exist. Off when power has just been applied to an enclosure, or when an enclosure fault exists.

Table 3-4. Status LEDs and Their Functions (page 2 of 3)

Location	LED Name	Color	Function
	Enclosure Status	Amber	Flashes when EMU is operational and performing locate. On when EMU is operational, but an enclosure fault exists. Off when EMU is operational, or power has just been applied to an enclosure, or when an EMU fault exists that is not an enclosure fault, or when an enclosure fault exists.
FC-AL I/O Module	Power-on	Middle Green	Lights when power is on and module is available for normal operation. If light is off, the module is nonoperational: check FCSAs, cables, and power supplies.
	Port 1	Bottom Green	Lights when carrier on Port 1 is operational.
	Port 2	Top Green	Lights when carrier on Port 2 is operational.
Fibre Channel ServerNet adapter (FCSA)	Power-on	Green	Lights when the adapter is receiving power.
	Service	Amber	Lights to indicate internal failure or service action required.
Gigabit Ethernet 4-port ServerNet adapter (G4SA)	Power-on	Green	Lights when the adapter is receiving power.
	Service	Amber	Lights to indicate internal failure or service action required.
LSU I/O PIC	Power-on	Green	Lights when power is on and adapter is available for normal operation.
	Service	Amber	Lights when a POST is in progress, board is being reset, or a fault exists.
LSU optics adapter connector	Power-on	Green	Lights when NonStop Blade Element optic or ServerNet link is functional.
LSU logic board	Power-on	Green	Lights when power is on and adapter is available for normal operation.
	Service	Amber	Lights when a POST is in progress, board is being reset, or a fault exists.

Table 3-4. Status LEDs and Their Functions (page 3 of 3)

Location	LED Name	Color	Function
NonStop Blade Element	Power-on	Flashing Green	Lights when power is on and Blade Element is available for normal operation.
		Flashing Yellow	Lights when Blade Element is in low power mode.
	Service	Steady Amber	Lights when a hardware or software fault exists.
P-switch PICs	Power-on	Flashing Blue	Lights when the system locator is activated.
		Green	Lights when power is on with PIC available for normal operation.
P-switch PIC ServerNet connector	Power-on	Amber	Lights when a fault exists.
		Green	Lights when a ServerNet link is functional.

Related Reading

For more information about monitoring, see the documentation listed in [Table 3-5](#).

Table 3-5. Related Reading for Monitoring

Task	Tool	For information, see...
Monitoring system hardware, including locating failed or failing FRUs	OSM Service Connection	OSM online help <i>OSM Service Connection User's Guide</i>
Using SCF, its commands and options, and device types and subtypes	SCF interface to subsystems	<i>SCF Reference Manual for H-Series RVUs</i> <i>SCF Reference Manual for the Storage Subsystem</i>
Monitoring clustered servers	OSM Service Connection	<i>ServerNet Cluster 6780 Operations Guide</i> <i>ServerNet Cluster Manual</i>

Monitoring EMS Event Messages

[When to Use This Section](#) on page 4-1

[What Is the Event Management Service \(EMS\)?](#) on page 4-1

[Tools for Monitoring EMS Event Messages](#) on page 4-1

[OSM Event Viewer](#) on page 4-2

[OSM Event Viewer](#) on page 4-2

[ViewPoint](#) on page 4-2

[Web ViewPoint](#) on page 4-2

[Related Reading](#) on page 4-2

When to Use This Section

Use this section for a brief description of the Event Management Service (EMS) and the tools used to monitor EMS event messages.

What Is the Event Management Service (EMS)?

The Event Management Service (EMS) is a collection of processes, tools, and interfaces that support the reporting and retrieval of event information. Information retrieved from EMS can help you to:

- Monitor your system or network environment
- Analyze circumstances that led up to a problem
- Detect failure patterns
- Adjust for changes in the run-time environment
- Recognize and handle critical problems
- Perform many other tasks required to maintain a productive computing operation

Tools for Monitoring EMS Event Messages

To view EMS event messages for an Integrity NonStop server, use one of these tools:

- OSM Event Viewer
- EMSDIST
- ViewPoint
- Web ViewPoint

OSM Event Viewer

The OSM Event Viewer is a browser-based event viewer. The OSM Event Viewer allows you to retrieve and view events from any EMS formatted log files (\$0, \$ZLOG, or an alternate collector) for rapid assessment of operating system problems.

To access the OSM Event Viewer, refer to [Launching OSM Applications](#) on page 1-11. For details on how to use the OSM Event Viewer, refer to the online help.

EMSDIST

The EMSDIST program is the object program for a printing, forwarding, or consumer distributor, any of which you can start with a TACL RUN command. This guide does not describe using EMSDIST. For more information, see the *Guardian User's Guide*.

ViewPoint

ViewPoint displays event messages about current or past events occurring anywhere in the network on a set of block-mode events screens. The messages can be errors, failures, warnings, and requests for operator actions. The events screens allow operators to monitor significant occurrences or problems in the network as they occur. Critical events or events requiring immediate action are highlighted.

Web ViewPoint

Web ViewPoint, a browser-based product, accesses the Event Viewer, Object Manager, and Performance Monitor subsystems. Web ViewPoint monitors and displays EMS events; identifies and lists all supported subsystems; manages NonStop server subsystems and user applications in a secure, automated, and customizable way; monitors and graphs performance attributes and trends; investigates and displays most active system processes; and offers simple navigation and a point-and-click command interface.

Related Reading

For more information about monitoring EMS event messages, see the documentation in [Table 4-1](#).

Table 4-1. Related Reading for Monitoring EMS Event Messages

Task	Tool	For information, see...
Viewing event logs	EMSDIST	<i>Guardian User's Guide</i>
	ViewPoint	<i>ViewPoint Manual</i>
	OSM Event Viewer	OSM Event Viewer online help

Processes: Monitoring and Recovery

[When to Use This Section](#) on page 5-1

[Types of Processes](#) on page 5-1

[System Processes on page 5-1](#)

[I/O Processes \(IOPs\) on page 5-2](#)

[Generic Processes on page 5-2](#)

[Monitoring Processes on page 5-3](#)

[Monitoring System Processes on page 3](#)

[Monitoring IOPs on page 4](#)

[Monitoring Generic Processes on page 4](#)

[Recovery Operations for Processes on page 5-6](#)

[Related Reading on page 5-6](#)

When to Use This Section

This section provides basic information about the different types of processes for Integrity NonStop servers. It gives a brief example of monitoring each type of process and provides information about the commands available for recovery operations.

Types of Processes

Three types of processes are of major concern to a system operator of an Integrity NonStop NS-series server:

- System processes
- I/O processes (IOPs)
- Generic processes

System Processes

A system process is a privileged process that is created during system load and exists continuously for a given configuration for as long as the processor remains operable. Examples of system processes include the memory manager, the monitor, and the I/O control processes.

I/O Processes (IOPs)

An I/O process (IOP) is a system process that manages communications between a processor and I/O devices. IOPs are often configured as fault-tolerant process pairs, and they typically control one or more I/O devices or communications lines. Each IOP is configured in a maximum of two processors, typically a primary processor and a backup processor.

An IOP provides an application program interface (API) that allows access to an I/O interface. A wide area network (WAN) communications line is an example of an I/O interface. IOPs configured using the SCF interface to the WAN subsystem manage the input and output functions for the ServerNet wide area network (SWAN) concentrator. Examples of IOPs include, but are not limited to, line-handler processes for Expand and other communications subsystems.

Generic Processes

Generic processes are configured by the SCF interface to the Kernel subsystem. They can be configured in one or more processors. Although sometimes called system-managed processes, generic processes can be either system processes or user-created processes. Any process that can be started from a TACL prompt can be configured as a generic process. Generic processes can be configured to have persistence; that is, to automatically restart if stopped abnormally.

Examples of generic processes:

- The \$ZZKRN Kernel subsystem manager process
- Other generic processes controlled by \$ZZKRN; for example:
 - The \$ZZSTO storage subsystem manager process
 - The \$ZZWAN wide area network (WAN) subsystem manager process
 - QIO processes
 - OSM server processes
 - The \$ZZLAN ServerNet LAN Systems Access (SLSA) subsystem manager process
 - The \$FCSMON fibre channel storage monitor

For more information, refer to the *SCF Reference Manual for the Kernel Subsystem*.

Monitoring Processes

This subsection briefly provides examples of some of the tools available to monitor processes. For some processes, such as IOPs, monitoring is more fully discussed in other manuals. In general, use this method to monitor processes:

1. Develop a list of processes that are crucial to the operation of your system.
2. Determine how each of these processes is configured.
3. Use the appropriate tool to monitor the process.

Monitoring System Processes

Check that the system processes are up and running in the processors as you intended. At a TACL prompt:

```
> STATUS *
```

This example shows partial output produced by the TACL STATUS * command:

Process	Pri	PFR	%WT	Userid	Program file	Hometerm
\$SYSTEM STARTUP 2> status *						
	0,0	201 P R	000	255,255	\$SYSTEM.SYS14.NMONTOR	\$YMIOP.#CLCI
	0,1	210 P	040	255,255	\$SYSTEM.SYS14.NMEMMAN	\$YMIOP.#CLCI
	0,2	210 P	051	255,255	\$SYSTEM.SYS14.NMSGERR	\$YMIOP.#CLCI
\$0	0,3	201 P	011	255,255	\$SYSTEM.SYS14.OPCOLL	\$YMIOP.#CLCI
	0,4	211 P	017	255,255	\$SYSTEM.SYS14.TMFMON	\$YMIOP.#CLCI
\$YMIOP	0,5	205 P	251	255,255	\$SYSTEM.SYS14.TMIOP	\$YMIOP.#CLCI
\$ZNUP	0,6	200 P	015	255,255	\$SYSTEM.SYS14.NZNUP	\$YMIOP.#CLCI
\$Z0	0,7	200 P	015	255,255	\$SYSTEM.SYS14.OCDIST	\$YMIOP.#CLCI
\$ZOPR	0,8	201 P	011	255,255	\$SYSTEM.SYS14.OAUX	\$YMIOP.#CLCI
\$ZCNF	0,9	200 P	001	255,255	\$SYSTEM.SYS14.TZCNF	\$YMIOP.#CLCI
\$ZTM00	0,11	200 P	017	255,255	\$SYSTEM.SYS14.TMFMON2	\$YMIOP.#CLCI
\$TMP	0,12	204 P	005	255,255	\$SYSTEM.SYS14.TMFTMP	\$YMIOP.#CLCI
\$ZL00	0,13	200 P	001	255,255	\$SYSTEM.SYS14.ROUT	\$ZHOME
\$NCP	0,14	199 P	011	255,255	\$SYSTEM.SYS14.NCPOBJ	\$ZHOME
\$ZEXP	0,15	150 P	001	255,255	\$SYSTEM.SYS14.OZEXP	\$ZHOME
\$CLCI	0,34	199	000	0,0	\$SYSTEM.SYS14.TACL	\$YMIOP.#CLCI
\$TRAK	0,40	146	000	255,255	\$SYSTEM.SYSTOOLS.QATRACK	\$ZHOME
\$Z00Y	0,43	150	015	255,255	\$SYSTEM.SYS14.FDIST	\$ZHOME
\$NULL	B 0,45	147	001	255,255	\$SYSTEM.SYSTEM.NULL	\$Z01J
\$ZNET	0,64	175 P	011	255,255	\$SYSTEM.SYS14.SCP	\$ZHOME
\$Z1RL	0,249	148	R 000	98,98	\$SYSTEM.SYS14.TACL	\$ZTNT.#PTBY5D
\$SYSTEM	0,257	220 P	317	255,255	\$SYSTEM.SYS14.TSYSDP2	\$YMIOP.#CLCI
\$ZHOME	0,292	199 P	001	255,255	\$SYSTEM.SYS14.ZHOME	\$YMIOP.#CLCI
\$ZM00	0,294	201 P	015	255,255	\$SYSTEM.SYS14.QIOMON	\$ZHOME
\$ZZWAN	0,295	180	011	255,255	\$SYSTEM.SYS14.WANMGR	\$ZHOME
\$ZZSTO	0,296	180 P	011	255,255	\$SYSTEM.SYS14.TZSTO	\$ZHOME
\$ZZLAN	0,297	199 P	015	255,255	\$SYSTEM.SYS14.LANMAN	\$ZHOME
\$ZZKRN	0,298	180 P	011	255,255	\$SYSTEM.SYS14.OZKRN	\$ZHOME
\$Z000	0,299	180 P	011	255,255	\$SYSTEM.SYS14.TZSTOSRV	\$ZHOME
\$ZLM00	0,300	200 P	015	255,255	\$SYSTEM.SYS14.LANMON	\$ZHOME
\$IXPOHO	0,301	199 P	355	255,255	\$SYSTEM.SYS14.LHOBJ	\$ZHOME
\$ZTXAE	0,330	145	015	255,255	\$SYSTEM.SYS14.SNMPTMUX	\$ZHOME
\$ZWBAF	0,333	179 P	015	255,255	\$SYSTEM.SYS14.WANBOOT	\$ZHOME
\$ZZW00	0,334	199 P	215	255,255	\$SYSTEM.SYS14.CONMGR	\$ZHOME
\$DSMSCM	0,335	220 P	317	255,255	\$SYSTEM.SYS14.TSYSDP2	\$ZHOME
\$DATA2	0,336	220 P	317	255,255	\$SYSTEM.SYS14.TSYSDP2	\$ZHOME
\$ZLOG	0,340	150	011	255,255	\$SYSTEM.SYS14.EMSACOLL	\$ZHOME
\$ZTH00	0,343	148 P	005	255,255	\$SYSTEM.SYS14.TFDSHLP	\$YMIOP.#CLCI
\$DSMSCM	0,344	220 P	317	255,255	\$SYSTEM.SYS14.TSYSDP2	\$ZHOME
\$Z1RM	1,80	148	005	255,255	\$SYSTEM.SYS14.TACL	\$ZTNT.#PTBY5D
\$ZPP01	1,280	160 P	015	255,255	\$SYSTEM.SYS14.OSSPS	\$YMIOP.#CLCI

\$ZLM01		1,342	200 P	015	255,255	\$SYSTEM.SYS14.LANMON	\$ZHOME
\$ZTC0	B	1,352	200 P	011	255,255	\$SYSTEM.SYS14.TCPIP	\$ZHOME
\$ZTNT	B	1,355	149	001	255,255	\$SYSTEM.SYS14.TELSERV	\$ZHOME
\$ZPORT	B	1,357	149	001	255,255	\$SYSTEM.SYS14.LISTNER	\$ZHOME
\$KLA9E		1,424	147	001	255,255	\$DATA2.KMZTT.LOGGER	\$ZTNT.#PTBY5D
\$ZTM02		2,5	200 P	017	255,255	\$SYSTEM.SYS14.TMFMON2	\$YMIOP.#CLCI
\$GRD2		2,243	147 P	001	255,255	\$DATA2.QA9050.RUNNER	\$ZTNT.#PTBY5CV
\$ZP02A	B	2,300	195	001	255,255	\$SYSTEM.ZRPC.PORTMAP	\$ZHOME
\$ZCMOM	B	2,303	150	001	255,255	\$SYSTEM.SYS14.CIMOM	\$ZHOME

Monitoring IOPs

For a list of manuals that provide more information about monitoring I/O processes (IOPs), refer to the *WAN Subsystem Configuration and Management Manual*, the *SWAN Concentrator and WAN Subsystem Troubleshooting Guide*, and the *Expand Configuration and Management Manual*.

Monitoring Generic Processes

Because generic processes are configured using the SCF interface to the Kernel subsystem, you specify the \$ZZKRN Kernel subsystem manager process when monitoring a generic process. These SCF commands are available for monitoring \$ZZKRN and other generic processes:

INFO	Displays configuration information for the specified objects
NAMES	Displays a list of subordinate object types and names for the specified objects
STATUS	Displays current status information about the specified objects

Monitoring the Status of \$ZZKRN

To monitor the status of the \$ZZKRN Kernel subsystem manager process, at a TACL prompt:

```
> SCF STATUS SUBSYS $ZZKRN
```

This example shows the output produced by this command:

1 -> STATUS SUBSYS \$ZZKRN		
NONSTOP KERNEL - Status SUBSYS \COMM.\$ZZKRN		
Name	State	Processes (conf/strd)
\COMM.\$ZZKRN	STARTED	(25/22)

Monitoring the Status of All Generic Processes

To monitor the status of all generic processes controlled by \$ZZKRN, at a TACL prompt:

```
> SCF STATUS PROCESS $ZZKRN.#*
```

This example shows the output produced by this command:

```

1-> STATUS PROCESS $ZZKRN.#*
NONSTOP KERNEL - Status PROCESS \DRP25.$ZZKRN.#CLCI-TACL

Symbolic Name      Name      State      Sub Primary Backup  Owner
                   PID       PID       PID       PID       ID
CLCI-TACL          $CLCI    STOPPED   None      None      None
MSGMON             $ZIM00   STARTED   0 ,306    None      255,255
MSGMON             $ZIM01   STARTED   1 ,291    None      255,255
MSGMON             $ZIM02   STARTED   2 ,285    None      255,255
MSGMON             $ZIM03   STARTED   3 ,280    None      255,255
MSGMON             $ZIM04   STARTED   4 ,280    None      255,255
MSGMON             $ZIM05   STARTED   5 ,280    None      255,255
MSGMON             $ZIM06   STARTED   6 ,280    None      255,255
MSGMON             $ZIM07   STARTED   7 ,280    None      255,255
MSGMON             $ZIM08   STARTED   8 ,280    None      255,255
MSGMON             $ZIM09   STARTED   9 ,280    None      255,255
MSGMON             $ZIM10   STARTED  10 ,280    None      255,255
MSGMON             $ZIM11   STOPPED   None      None      None
MSGMON             $ZIM12   STOPPED   None      None      None
MSGMON             $ZIM13   STOPPED   None      None      None
MSGMON             $ZIM14   STOPPED   None      None      None
MSGMON             $ZIM15   STOPPED   None      None      None
OSM-APPSRV        $ZOSM    STARTED   2 ,292    None      255,255
OSM-CIMOM          $ZCMOM   STARTED   2 ,294    3 ,288    255,255
OSM-CONFLH-RD     $ZOLHI   STOPPED   None      None      None
OSM-OEV            $ZOEV    STARTED   2 ,290    None      255,255
QATRAK             $TRAK    STARTED   0 ,17     None      255,255
QIOMON             $ZM00    STARTED   0 ,290    None      255,255
QIOMON             $ZM01    STARTED   1 ,280    None      255,255
QIOMON             $ZM02    STARTED   2 ,280    None      255,255
QIOMON             $ZM03    STARTED   3 ,279    None      255,255
QIOMON             $ZM04    STARTED   4 ,279    None      255,255
QIOMON             $ZM05    STARTED   5 ,279    None      255,255
QIOMON             $ZM06    STARTED   6 ,279    None      255,255
QIOMON             $ZM07    STARTED   7 ,279    None      255,255
QIOMON             $ZM08    STARTED   8 ,279    None      255,255
QIOMON             $ZM09    STARTED   9 ,279    None      255,255
QIOMON             $ZM10    STARTED  10 ,279    None      255,255
QIOMON             $ZM11    STOPPED   None      None      None
QIOMON             $ZM12    STOPPED   None      None      None
QIOMON             $ZM13    STOPPED   None      None      None
QIOMON             $ZM14    STOPPED   None      None      None
QIOMON             $ZM15    STOPPED   None      None      None
RTACL              $RTACL   STOPPED   None      None      None
SCP                $ZNET    STARTED   0 ,14     1 ,13     255,255
SP-EVENT           $ZSPE    STARTED   0 ,309    None      255,255
TFDSHLP           $ZTH00   STARTED   0 ,310    None      255,255
TFDSHLP           $ZTH01   STARTED   1 ,292    None      255,255
TFDSHLP           $ZTH02   STARTED   2 ,286    None      255,255
TFDSHLP           $ZTH03   STARTED   3 ,281    None      255,255
TFDSHLP           $ZTH04   STARTED   4 ,281    None      255,255
TFDSHLP           $ZTH05   STARTED   5 ,281    None      255,255
TFDSHLP           $ZTH06   STARTED   6 ,281    None      255,255
TFDSHLP           $ZTH07   STARTED   7 ,281    None      255,255
TFDSHLP           $ZTH08   STARTED   8 ,281    None      255,255
TFDSHLP           $ZTH09   STARTED   9 ,281    None      255,255
TFDSHLP           $ZTH10   STARTED  10 ,281    None      255,255
TFDSHLP           $ZTH11   STOPPED   None      None      None
TFDSHLP           $ZTH12   STOPPED   None      None      None

```

TFDSHLP	\$ZTH13	STOPPED	None	None	
TFDSHLP	\$ZTH14	STOPPED	None	None	
TFDSHLP	\$ZTH15	STOPPED	None	None	
ZEXP	\$ZEXP	STARTED	0 ,13	1 ,15	255 ,255
ZHOME	\$ZHOME	STARTED	0 ,289	1 ,295	255 ,255
ZLOG	\$ZLOG	STARTED	0 ,308	1 ,329	255 ,255
ZZKRN	\$ZZKRN	STARTED	0 ,293	1 ,319	255 ,255
ZZLAN	\$ZZLAN	STARTED	0 ,292	1 ,297	255 ,255
ZZSCL	\$ZZSCL	STARTED	1 ,290	2 ,279	255 ,255
ZZSMN	\$ZZSMN	STARTED	1 ,289	2 ,282	255 ,255
ZZSTO	\$ZZSTO	STARTED	0 ,291	1 ,320	255 ,255
ZZWAN	\$ZZWAN	STARTED	2 ,296	3 ,289	255 ,255

In nearly all circumstances, items that are essential to system operations that must be running at all times restart automatically if they are stopped for any reason while the NonStop Kernel operating system is running.

Some OSM processes stop after executing a macro that runs during system load or during the reload of processor 0 or 1. Those processes include \$ZOLHI.

Optionally, you can also configure other processes such as the Expand subsystem manager process, \$ZEXP, and the Safeguard monitor process, \$ZSMP, as generic processes.

Recovery Operations for Processes

For recovery operations on generic processes, use the SCF interface to the Kernel subsystem and specify the PROCESS object. These SCF commands are available for controlling generic processes:

- | | |
|-------|---|
| ABORT | Terminates operation of a generic process. This command is not supported for the subsystem manager processes. |
| START | Initiates the operation of a generic process. |

Generic processes that are configured to be persistent usually do not require operator intervention for recovery. In most circumstances, persistent generic processes restart automatically.

For recovery operations on IOPs, refer to the *WAN Subsystem Configuration and Management Manual*, the *SWAN Concentrator and WAN Subsystem Troubleshooting Guide*, and the *Expand Configuration and Management Manual*.

For recovery operations on system processes, refer to the *Guardian User's Guide*.

Related Reading

For more information about generic processes and the SCF interface to the Kernel subsystem, refer to the *SCF Reference Manual for the Kernel Subsystem*.

For more information about IOPs, refer to the *WAN Subsystem Configuration and Management Manual*, the *SWAN Concentrator and WAN Subsystem Troubleshooting Guide*, and the *Expand Configuration and Management Manual*.

6

Communications Subsystems: Monitoring and Recovery

[When to Use This Section](#) on page 6-1

[Communications Subsystems](#) on page 6-1

[Local Area Networks \(LANs\) and Wide Area Networks \(WANs\)](#) on page 6-2

[Monitoring Communications Subsystems and Their Objects](#) on page 6-4

[Monitoring the SLSA Subsystem](#) on page 6-4

[Monitoring the WAN Subsystem](#) on page 6-6

[Monitoring the NonStop TCP/IP Subsystem](#) on page 6-9

[Monitoring Line-Handler Process Status](#) on page 6-10

[Tracing a Communications Line](#) on page 6-12

[Recovery Operations for Communications Subsystems](#) on page 6-13

[Related Reading](#) on page 6-13

When to Use This Section

Use this section to determine where to find more information about monitoring and recovery operations for communications devices such as ServerNet adapters, printers, and spoolers; communications lines; and communications processes such as WAN IOPs.

Communications Subsystems

The software that provides users of Integrity NonStop systems with access to a set of communications services is called a communications subsystem. Because connectivity is an important part of online transaction processing (OLTP), HP offers a variety of communications products that support a wide range of applications.

Communication between specific devices or networks is typically achieved using several communications products or subsystems. These products are related as components in a layered structure. To accomplish the required connection, higher-level components—for example, NonStop TCP/IP processes—use the services of lower-level components such as the ServerNet LAN Systems Access (SLSA) subsystem.

The same higher-level component can often use any of several lower-level components; thus, the Expand subsystem—which consists of multiple processes on a node—can use the NonStop TCP/IP subsystem, the X.25 Access Method (X.25 AM),

or other communication interface options to provide data transmissions over local area networks (LANs) or wide area networks (WANs), respectively. Similarly, multiple higher-level components can use the services of a single lower-level component.

Local Area Networks (LANs) and Wide Area Networks (WANs)

Two important communications interfaces for LANs and WANs on Integrity NonStop servers are the SLSA subsystem and the WAN subsystem.

The SLSA subsystem supports parallel LAN I/O operations, allowing Integrity NonStop NS-series servers to communicate across the ServerNet fabrics and access Ethernet devices through various LAN protocols. SLSA also communicates with the appropriate adapter type over the ServerNet fabrics. Adapters supported on Integrity NonStop systems include:

- Gigabit Ethernet 4-port adapter (G4SA)
- Fibre Channel ServerNet adapter (FCSA) (for the Storage subsystem)

I/O adapter module (IOAM) enclosures enable I/O operations to take place between Integrity NonStop servers and some Fibre Channel storage devices. See the *Modular I/O Installation and Configuration Guide* for more information.

Adapters supported on NonStop S-series servers that can be accessed through Expand over IP, include:

- ATM 3 ServerNet adapter (ATM3SA)
- Ethernet 4 ServerNet adapter (E4SA)
- Fast Ethernet ServerNet adapter (FESA)
- Gigabit Ethernet ServerNet adapter (GESA)
- Gigabit Ethernet 4-Port ServerNet adapter (G4SA)
- Multifunction I/O board (MFIOB) in the processor multifunction (PMF) customer-replaceable unit (CRU) and I/O multifunction (IOMF) CRU
- Token-Ring ServerNet adapter (TRSA)

For further information, refer to the *Introduction to Networking for NonStop NS-Series Servers*.

In addition to the adapters, the SLSA subsystem supports these objects:

- Processes
- Monitors
- ServerNet addressable controllers (SACs)
- Logical interfaces (LIFs)
- Filters

- Physical interfaces (PIFs)

Processes that use the SLSA subsystem to send and receive data on a LAN attached to an Integrity NonStop server are called LAN service providers. Two service providers—the NonStop TCP/IP and NonStop TCP/IPv6 subsystems and the Port Access Method (PAM)—are currently supported. They provide access for these subsystems:

LAN Service Provider	Subsystems Supported
NonStop TCP/IP subsystem, NonStop TCP/IPv6 subsystem	The Expand subsystem, which provides Expand-over-IP connections.
Port Access Method (PAM)	Ethernet and token-ring LANs. The OSI/AS, OSI/TS, SNAX/XF, and SNAX/APN subsystems communicate with SLSA through the PAM subsystem.

Processes, user applications, and subsystems that use the SLSA subsystem and related LAN providers to connect to an FCSA or G4SA attached to an Integrity NonStop NS-series server are called LAN clients. For example, the WAN subsystem is a client of the SLSA subsystem because the SLSA subsystem provides the WAN subsystem access to the ServerNet wide area network (SWAN) concentrator through the LAN.

The WAN subsystem is used to control access to the SWAN concentrator. Depending on your configuration, it can be used to configure and manage both WAN and LAN connectivity for these communication subsystem objects:

Object	Connectivity By
AM3270	Line-handler processes
Asynchronous Terminal Process 6100 (ATP6100)	Line-handler processes
Communications Process subsystem (CP6100)	Line-handler processes
EnvoyACP/XF	Line-handler processes
Envoy subsystem	Line-handler processes
Expand	Subsystem network control process and line-handler processes
ServerNet cluster (Expand-over-ServerNet)	Line-handler processes
SNAX/APN	Subsystem service manager process and line-handler processes
SNAX/XF	Subsystem service manager process and line-handler processes
TR3271	Line-handler processes
X25AM	Line-handler processes

You can define these communications subsystem objects as WAN subsystem devices.

Monitoring Communications Subsystems and Their Objects

Monitoring and recovery operations for communications subsystems can be complex. An error in any of the components—service providers, clients, objects, adapters, processes, and so on—can generate multiple error messages from many interdependent subsystems and processes. Analyzing and solving an error that originates in an object controlled by a LAN or a WAN often requires that you methodically gather status information about the affected services and then eliminate objects that are working normally.

Detailed monitoring and recovery techniques for devices and processes related to communications subsystems are discussed in detail in the manuals for each subsystem. For more information, refer to [Related Reading](#) on page 6-13.

This guide provides some basic commands you can use to identify and resolve common problems. Your most powerful tool for monitoring and collecting information about subsystem objects is the SCF facility. You can use SCF commands to get information and status for subsystem objects by name, device type, or device subtype.

Subdevices are defined if a subsystem potentially operates on numerous, separately addressable objects, such as stations on a multipoint line; the line is a device, and the stations are subdevices.

For a list of subsystems with their device type numbers and device subtypes, see [Using SCF to Determine Your System Configuration](#) on page 2-5.

Monitoring the SLSA Subsystem

This subsection describes how to obtain the status of adapters, SACs, LIFs, and PIFs. For more information on the SLSA subsystem, refer to the *LAN Configuration and Management Manual*.

Monitoring the Status of an Adapter and Its Components

1. To monitor the status of an adapter:

```
> SCF STATUS ADAPTER adapter-name
```

A listing similar to this example is sent to your home terminal:

```

->STATUS ADAPTER $ZZLAN.G11123

SLSA Status ADAPTER

Name                State
$ZZLAN.G11123      STARTED
```

This example shows the listing displayed when checking all adapters on \$ZZLAN:

> SCF STATUS ADAPTER \$ZZLAN.*

```

1->STATUS ADAPTER $ZZLAN.*

SLSA Status ADAPTER

Name                State
$ZZLAN.G11121       STARTED
$ZZLAN.G11122       STARTED
$ZZLAN.G11123       STARTED
$ZZLAN.G11124       STARTED
$ZZLAN.G11125       STARTED
$ZZLAN.MIOE0        STARTED
$ZZLAN.MIOE1        STARTED
    
```

- The SAC object corresponds directly to the hardware on an adapter. A SAC is a component of an adapter and can support one or more PIFs. To monitor the status of a SAC:

> SCF STATUS SAC *sac-name*

A listing similar to this example is sent to your home terminal:

```

1->STATUS SAC $ZZLAN.G11123.0

SLSA Status SAC

Name                Owner    State    Trace Status
$ZZLAN.G11123.0    1        STARTED  ON
    
```

This example shows a listing of the status of all SACs on \$ZZLAN.G11123:

> SCF STATUS SAC \$ZZLAN.G11123*

```

->STATUS SAC $ZZLAN.G11123*

SLSA Status SAC

Name                Owner    State    Trace Status
$ZZLAN.G11123.0    1        STARTED  ON
    
```

- The PIF object corresponds directly to hardware on the adapter. A PIF is the physical connection to the LAN. To monitor the status of a PIF:

> SCF STATUS PIF *pif-name*

A listing similar to this example is sent to your home terminal:

```

->STATUS PIF $ZZLAN.G11123.0

SLSA Status PIF

Name                State    Trace Status
$ZZLAN.G11123.0.A  STARTED  ON
    
```

This example shows a listing of the status of all PIFs on \$ZZLAN.G11123:

```
> SCF STATUS PIF $ZZLAN.G11123.*
```

```
->STATUS PIF $ZZLAN.G11123.*

SLSA Status PIF

Name                State      Trace Status
$ZZLAN.G11123.0.A   STARTED   ON
$ZZLAN.G11123.0.B   STARTED   ON
$ZZLAN.G11123.0.C   STOPPED   OFF
$ZZLAN.G11123.0.D   STARTED   ON
```

4. The LIF provides an interface to the PIF. The LIF object corresponds to logical processes that handle data transferred between the LAN and a system using the ServerNet architecture. To monitor the status of a LIF:

```
> SCF STATUS LIF lif-name
```

A listing similar to this example is sent to your home terminal:

```
->STATUS LIF $ZZLAN.L11021A

SLSA Status LIF

Name                State      Access State
$ZZLAN.L11021A      STARTED    UP
```

This example shows a detailed listing of the status of the LIF on \$ZZLAN.L11021A:

```
> SCF STATUS LIF $ZZLAN.L11021A , DETAIL
```

```
->STATUS LIF $ZZLAN.L11021A , DETAIL

SLSA Detailed Status LIF \SYS.$ZZLAN.L11021A

Access State..... UP
CPUs with Data Path..... ( 0, 1, 2 )
Potential Access CPUs.... ( 0, 1, 2, 3 )
State..... STARTED
Trace Filename.....
Trace Status.....
```

Monitoring the WAN Subsystem

This subsection describes how to obtain the status of SWAN concentrators, data communications devices, processes, and CLIPs. For more information on the WAN subsystem, see the *WAN Subsystem Configuration and Management Manual*.

Monitoring Status for a SWAN Concentrator

To display the current status for a SWAN concentrator:

```
> SCF STATUS ADAPTER $ZZWAN.#concentrator-name
```

The system displays a listing similar to:

```
-> status adapter $zzwan.#s01

WAN Manager STATUS ADAPTER for ADAPTER \TAHITI.$ZZWAN.#S01
  State..... STARTED

Number of clips. 3

Clip 1 status : CONFIGURED
Clip 2 status : CONFIGURED
Clip 3 status : CONFIGURED
```

To display the status for all SWAN concentrators configured for your system:

```
> SCF STATUS ADAPTER $ZZWAN.*
```

The system displays a listing similar to:

```
1-> STATUS ADAPTER $ZZWAN.*

WAN Manager STATUS ADAPTER for ADAPTER \COMM.$ZZWAN.#SWAN1
  State..... STARTED

Number of clips. 3

Clip 1 status : CONFIGURED
Clip 2 status : CONFIGURED
Clip 3 status : CONFIGURED

WAN Manager STATUS ADAPTER for ADAPTER \COMM.$ZZWAN.#SWAN2
  State..... STARTED

Number of clips. 3

Clip 1 status : CONFIGURED
Clip 2 status : CONFIGURED
Clip 3 status : CONFIGURED
```

Monitoring Status for a Data Communications Device

To verify that a WAN subsystem device is in the STARTED state:

```
> SCF STATUS DEVICE $ZZWAN.#device-name
```

The system displays a listing similar to:

```
-> status DEVICE $zzwan.#IP01

WAN Manager STATUS DEVICE for DEVICE \COWBOY.$ZZWAN.#IP01
  STATE .....STARTED

LDEV number....173

PPIN.....2, 13  BPIN.....3, 11
```

Monitoring WAN Processes

To display the status of all WAN subsystem processes—configuration managers, TCP/IP processes, WANBoot processes:

```
> SCF STATUS PROCESS $ZZWAN.*
```

The system displays a listing similar to:

```
-> STATUS PROCESS $ZZWAN.*  
  
WAN Manager STATUS PROCESS for PROCESS \COMM.$ZZWAN.#5  
State :..... STARTED  
  
LDEV Number..... 66  
PPIN..... 5 ,264 Process traced.. NO  
  
WAN Manager STATUS PROCESS for PROCESS \COMM.$ZZWAN.#4  
State :..... STARTED  
  
LDEV Number..... 67  
PPIN..... 4 ,264 Process traced.. NO  
  
WAN Manager STATUS PROCESS for PROCESS \COMM.$ZZWAN.#ZTF00  
State :..... STARTED  
  
PPIN..... 4 ,342  
  
WAN Manager STATUS PROCESS for PROCESS \COMM.$ZZWAN.#SWB1  
State :..... STARTED  
  
PPIN..... 4 ,275 BPIN..... 5 ,302  
  
WAN Manager STATUS PROCESS for PROCESS \COMM.$ZZWAN.#ZTF01  
State :..... STARTED  
  
PPIN..... 5 ,340  
  
WAN Manager STATUS PROCESS for PROCESS \COMM.$ZZWAN.#SWB0  
State :..... STARTED  
  
PPIN..... 4 ,274 BPIN..... 5 ,303
```

To monitor a single WANBoot process, type:

```
> SCF STATUS PROCESS $ZZWAN.#boot-process
```

The system displays a listing similar to:

```
-> status PROCESS $ZZWAN.#ZB017  
  
WAN Manager STATUS PROCESS for PROCESS \ICEBAT.$ZZWAN.#ZB017  
STATE:.....STARTED  
  
PPIN.....0 ,278 BPIN.....0 ,282
```

Monitoring CLIPs

To display the current status for a CLIP:

```
> SCF STATUS SERVER $ZZWAN.#concentrator-name.clip-num
```

Values for the CLIP number are 1, 2, or 3.

The system displays a listing similar to:

```
-> status server $zzwan.#s01.1

WAN Manager STATUS SERVER for CLIP \COWBOY.$ZZWAN.#S01.1
STATE :.....STARTED

PATH A.....: CONFIUGRED
PATH B.....: CONFIGURED

NUMBER of lines. 2

Line.....0      : $SAT23A
Line.....1      : $SAT23B
```

Monitoring the NonStop TCP/IP Subsystem

This subsection describes how to obtain the status for NonStop TCP/IP processes, routes, and subnets. For additional information, refer to the *TCP/IP Configuration and Management Manual*. For NonStop TCP/IPv6, refer to the *TCP/IPv6 Configuration and Management Manual*.

Monitoring the NonStop TCP/IP Process

To display the dynamic state of a NonStop TCP/IP process, first list the names of all NonStop TCP/IP processes:

```
-> SCF LISTDEV TCPIP
```

Then type:

```
> SCF STATUS PROCESS tcp/ip-process-name
```

where *tcp/ip-process-name* is the name of the process you want information about.

The system displays a listing similar to this output, which is for process \$ZTCO:

```
-> Status Process $ZTCO

TCPIP Status PROCESS \SYSA.$ZTCO

Status: STARTED
PPID.....( 0,107)          BPID..... ( 1. 98)

Proto State      Laddr      Lport      Faddr      Fport      SendQ      RecvQ
TCP    TIME-WAIT    130.252.12.3 ftp-data   130.252.12.152 11089      0         0
TCP    TIME-WAIT    130.252.12.3 ftp-data   130.252.12.152 63105      0         0
TCP    ESTAB        130.252.12.3 ftp        130.252.12.152 57441      0         0
TCP    TIME-WAIT    130.252.12.3 smtp      130.252.12.8   3309       0         0
```

Monitoring NonStop TCP/IP Routes

To display status information for all NonStop TCP/IP routes:

```
> SCF STATUS ROUTE $ZTCO.*
```

The system displays a listing similar to:

```
1-> Status Route $ZTCO.*
TCPIP Status ROUTE \SYSA.$ZTCO.*
Name                Status                RefCnt
#ROU11              STARTED                0
#ROU9                STARTED                0
#ROU12              STARTED                0
#ROU8                STARTED                1
#ROU3                STOPPED                0
```

Monitoring NonStop TCP/IP Subnets

To obtain the status of all NonStop TCP/IP subnets:

```
> SCF STATUS SUBNET $ZTCO.*
```

The system displays a listing similar to:

```
1-> STATUS SUBNET $ZTCO.*
TCPIP Status SUBNET \SYSA.$ZTCO.*
Name                Status
#LOOP0              STARTED
#EN1                 STARTED
```

Monitoring Line-Handler Process Status

A line-handler process is a component of a data communications subsystem. It is an I/O process that transmits and receives data on a communications line, either directly or by communicating with another I/O process. This subsection explains how to monitor the status of a line-handler process on your system or on another system in your network to which you have remote access.

To check the status of a line-handler process on your system:

```
> SCF STATUS LINE $line
```

A listing similar to this example is sent to your home terminal:

```
1-> STATUS LINE $LHPLIN1
EXPAND Status LINE
Name                State                PPID                BPID                ConMgr-LDEV
$LHCS6S             STARTED              1, 20              2,25               49
```

This listing shows that the Expand line-handler process being monitored is up and functioning normally.

The data shown in the report means:

Name	Specifies the name of the object
State	Indicates the summary state of the object, which is either STARTED, STARTING, DIAGNOSING (for SWAN concentrators only), or STOPPED
PPID	Specifies the primary process ID
BPID	Specifies the backup process ID
ConMgr-LDEV	Contains the LDEV of the concentrator manager process. This field applies only to SWAN concentrator lines.

If any state other than STARTED appears, check the meaning of the state in [SCF Object States](#) on page 3-14. Depending upon the type of problem, follow your established procedures for problem reporting and escalation.

Examples

To check the detailed status of line \$LHCS6S:

```
> SCF STATUS LINE $LHCS6S, DETAIL
```

A listing such as this output is sent to your home terminal:

```

-> STATUS LINE $LHCS6S, DETAIL

PPID..... ( 3, 24) BPID..... ( 2, 24)
State..... STOPPED Path LDEV..... 50
Trace Status..... OFF Clip Status..... UNLOADED
ConMgr-LDEV..... 49
Path-prim
Path-alter

```

To display the status of all the Expand lines that are currently active on your system, enter this INFO PROCESS command for the Expand manager process \$NCP:

```
-> INFO PROCESS $NCP, LINESET
```

The system displays a listing similar to this output. The NEIGHBOR field displays the system to which a given line connects, and the STATUS field indicates whether the line is up:

```

1-> INFO PROCESS $NCP, LINESET
EXPAND  Info  PROCESS      $NCP      , LINESET
        LINESETS AT \COMM      (116) #LINESETS=35 TIME:  JUL  9,2001 19:28:04
LINESET  NEIGHBOR  LDEV  TF  PID  LINE  LDEV  STATUS  FileErr#
  1  \CYCLONE (206)  363  200K ( 0, 287)  1  363  READY
  2  \SNAX (118)  353  200K ( 5, 333)  1  353  READY
  3  \TESS (194)  554  200K ( 8, 279)  1  554  READY
  4  \TSII (099)  556  200K ( 2, 265)  1  556  READY
  5  \ESP (163)  365  200K ( 1, 274)  1  365  READY
  6  \SVLDEV (077)  538  200K ( 7, 265)  1  538  READY
.
.
.
 27  \SIERRA (012)  183  10K ( 4, 290)  1  183  READY
 28  \PRUNE (175)  677  200K ( 5, 334)  1  677  READY
 29  \OPMAN (252)  276  790K ( 5, 294) NPT  1  276  READY
 30  \SOCIAL (045)  165  790K ( 8, 280)  1  165  READY
 31  \NCCORP2 (080)  295  790K ( 8, 264)  1  295  READY
 32  \CS8 (152)  323  -- -- -----  1  323 NOT READY (124)
 33  \CORE (241)  324  -- -- -----  1  324 NOT READY (124)
 34  \SUNTEC (062)  367  790K ( 5, 293) NPT  1  367  READY
 35  \CS8 (152)  368  -- -- -----  1  368 NOT READY (124)

```

Tracing a Communications Line

Use the SCF TRACE command to trace the operation of a communications line. The line continues normal operation while being traced, but it passes all its message traffic to a trace procedure. Tracing enables you to see the history of a communications line, including its internal processing.

You can display trace files by using the commands available in the PTrace program. For information about PTrace, refer to the *PTrace Reference Manual*. For information about configuring a trace by using the SCF TRACE command, refer to the configuration and management manual for the communications subsystem you want to trace.

Recovery Operations for Communications Subsystems

Some general troubleshooting guidelines are:

- Examine the contents of the event message log for the subsystem. For example, the WAN subsystem or Kernel subsystem might have issued an event message that provides information about the process failure. Event messages returned by the WAN subsystem and SWAN concentrator are described in the WANMGR and TRAPMUX sections of the *Operator Messages Manual*, respectively.
- HP provides a comprehensive library of troubleshooting guides for the communications subsystems. Attempt to analyze the problems and restart the process or object using the commands described in the appropriate manual listed in [Related Reading](#) on page 6-13. If you are unable to start a required process or object, contact your service provider.

Related Reading

For more information about monitoring and performing recovery operations for communications subsystems, see the manuals listed in [Table 6-1](#). The appropriate manual to use depends on how your system is configured.

For example, if a process is configured using the SCF interface to the WAN subsystem and then reconfigured with the SCF interface to another subsystem, only the SCF interface to the other subsystem would provide current information about the configuration. The SCF interface to the WAN subsystem would provide only information about the configuration before it changed.

Table 6-1. Related Reading for Communications Lines and Devices (page 1 of 2)

For Information About...	Refer to...
General information about communications subsystems	<i>Introduction to Networking for HP NonStop NS-Series Servers</i>
Using SCF to monitor generic processes	<i>SCF Reference Manual for the Kernel Subsystem</i>
Using SCF to monitor the SLSA subsystem as well as Ethernet addressable devices, such as ServerNet adapters	<i>LAN Configuration and Management Manual</i>

Table 6-1. Related Reading for Communications Lines and Devices (page 2 of 2)

For Information About...	Refer to...
Using SCF to monitor WAN communications lines for devices and intersystem communications protocols	<i>WAN Subsystem Configuration and Management Manual</i>
Using SCF to monitor a specific device or communications protocol product; troubleshooting specific communications subsystems and protocols	<i>Asynchronous Terminals and Printer Processes Configuration and Management Manual</i> <i>ATM Adapter Installation and Support Guide</i> <i>ATM Configuration and Management Manual</i> <i>CP6100 Configuration and Management Manual</i> <i>EnvoyACP/XF Configuration and Management Manual</i> <i>Expand Configuration and Management Manual</i> <i>Fibre Channel ServerNet Adapter Installation and Support Guide</i> <i>Gigabit Ethernet 4-Port Adapter Installation and Support Guide</i> <i>PAM Configuration and Management Manual</i> <i>QIO Configuration and Management Manual</i> <i>SCF Reference Manual for H-Series RVUs</i> <i>ServerNet Cluster Manual</i> <i>SNAX/XF and SNAX/APN Configuration and Management Manual</i> <i>SWAN Concentrator and WAN Subsystem Troubleshooting Guide</i> <i>TCP/IPv6 Configuration and Management Manual</i> <i>TCP/IP Configuration and Management Manual</i> <i>Token-Ring Adapter Installation and Support Guide</i> <i>X25AM Configuration and Management Manual</i>

7

ServerNet Resources: Monitoring and Recovery

[ServerNet Communications Network](#) on page 7-1

[System I/O ServerNet Connections](#) on page 7-4

[Monitoring the Status of the ServerNet Fabrics](#) on page 7-4

[Monitoring the ServerNet Fabrics Using OSM](#) on page 7-5

[Monitoring the ServerNet Fabrics Using SCF](#) on page 7-6

[Related Reading](#) on page 7-8

When to Use This Section

Use this section to learn about monitoring and performing recovery operations for the internal and external ServerNet fabrics, and to understand how and when an Integrity NonStop NS-series system can be connected to legacy NonStop S-series I/O enclosures.

Notes. Integrity NonStop NS16000 systems support connectivity to NonStop S-series I/O enclosures, Integrity NonStop NS14000 and NS1000 systems do not. For more information, see [Differences Between Integrity NonStop NS-Series Systems](#) on page 2-2.

An Integrity NonStop NS16000 system can be part of the same ServerNet cluster as NonStop S-series systems, an Integrity NonStop NS14000 system cannot be. For more information, see the *ServerNet Cluster Supplement for Integrity NonStop NS-Series Servers*.

Integrity NonStop NS1000 systems do not support ServerNet clusters.

All Integrity NonStop system I/O is performed through the ServerNet system area network (SAN). LSU logic boards connect the SAN to the replicated four-way microprocessors on Integrity NonStop systems (except for Integrity NonStop NS1000 systems, which have no LSUs; see [System I/O ServerNet Connections](#) on page 7-4).

ServerNet Communications Network

The ServerNet communications network is a high-speed network within an Integrity NonStop system that connects processors to each other and to peripheral controllers. This network offers the connectivity of a standard network, but it does not depend on shared resources such as interprocessor buses or I/O channels. Instead, the ServerNet communications network uses the ServerNet architecture, which is wormhole-routed, full-duplex, packet-switched, and point-to-point. This network offers low latency, low software overhead, high bandwidth, and parallel operation.

In the ServerNet architecture, each processor maintains two independent paths to other processors, I/O devices, and ServerNet adapters. These dual paths can be used

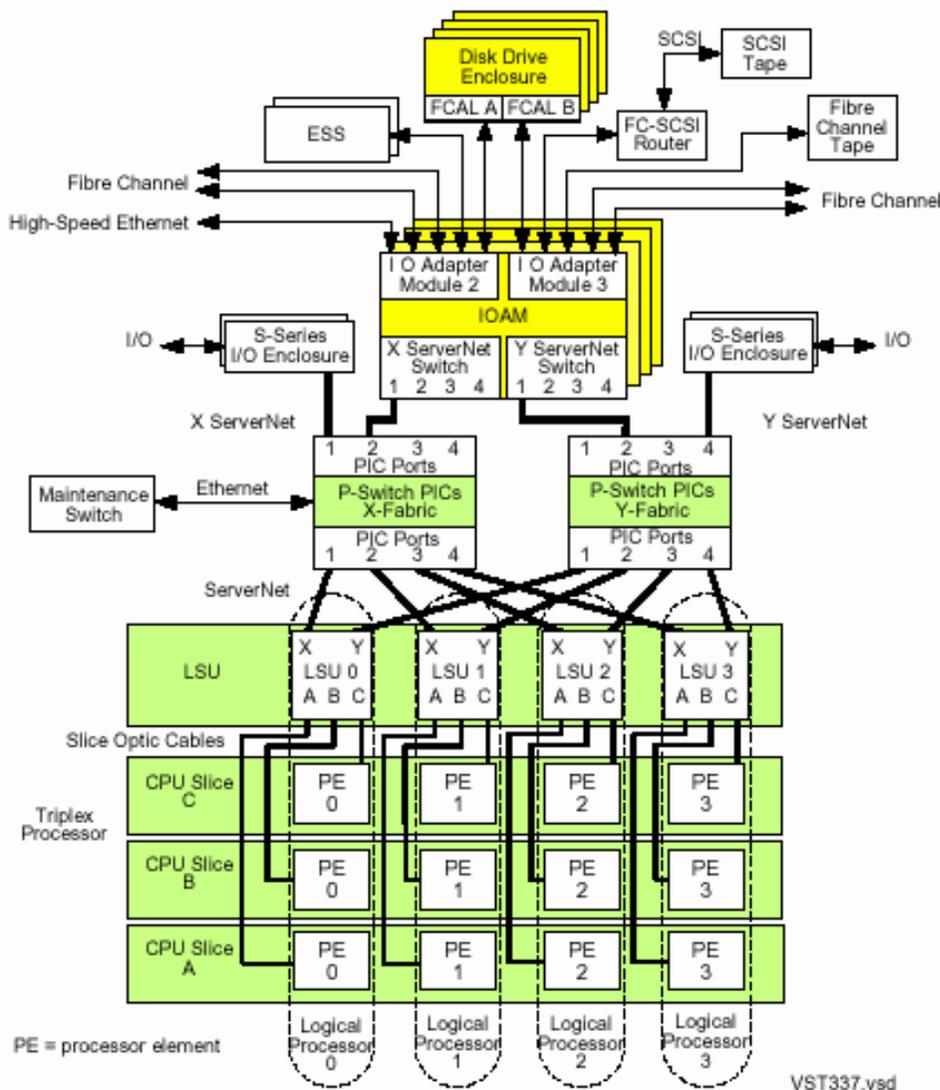
simultaneously to improve performance, and to ensure that no single failure disrupts communications among the remaining system components.

A ServerNet adapter provides the interface between a ServerNet fabric and the Fibre Channel and Ethernet links. A ServerNet adapter contains a ServerNet bus interface (SBI) and one or more ServerNet addressable controllers (SACs).

Integrity NonStop NS16000 ServerNet Connectivity

An Integrity NonStop NS16000 system uses the ServerNet fabric for interconnections between the LSUs, p-switches, and IOAMs, enabling an Integrity NonStop system to be connected to legacy NonStop S-series enclosures. [Figure 7-1](#) shows a logical representation of a complete system with the X and Y ServerNet fabrics.

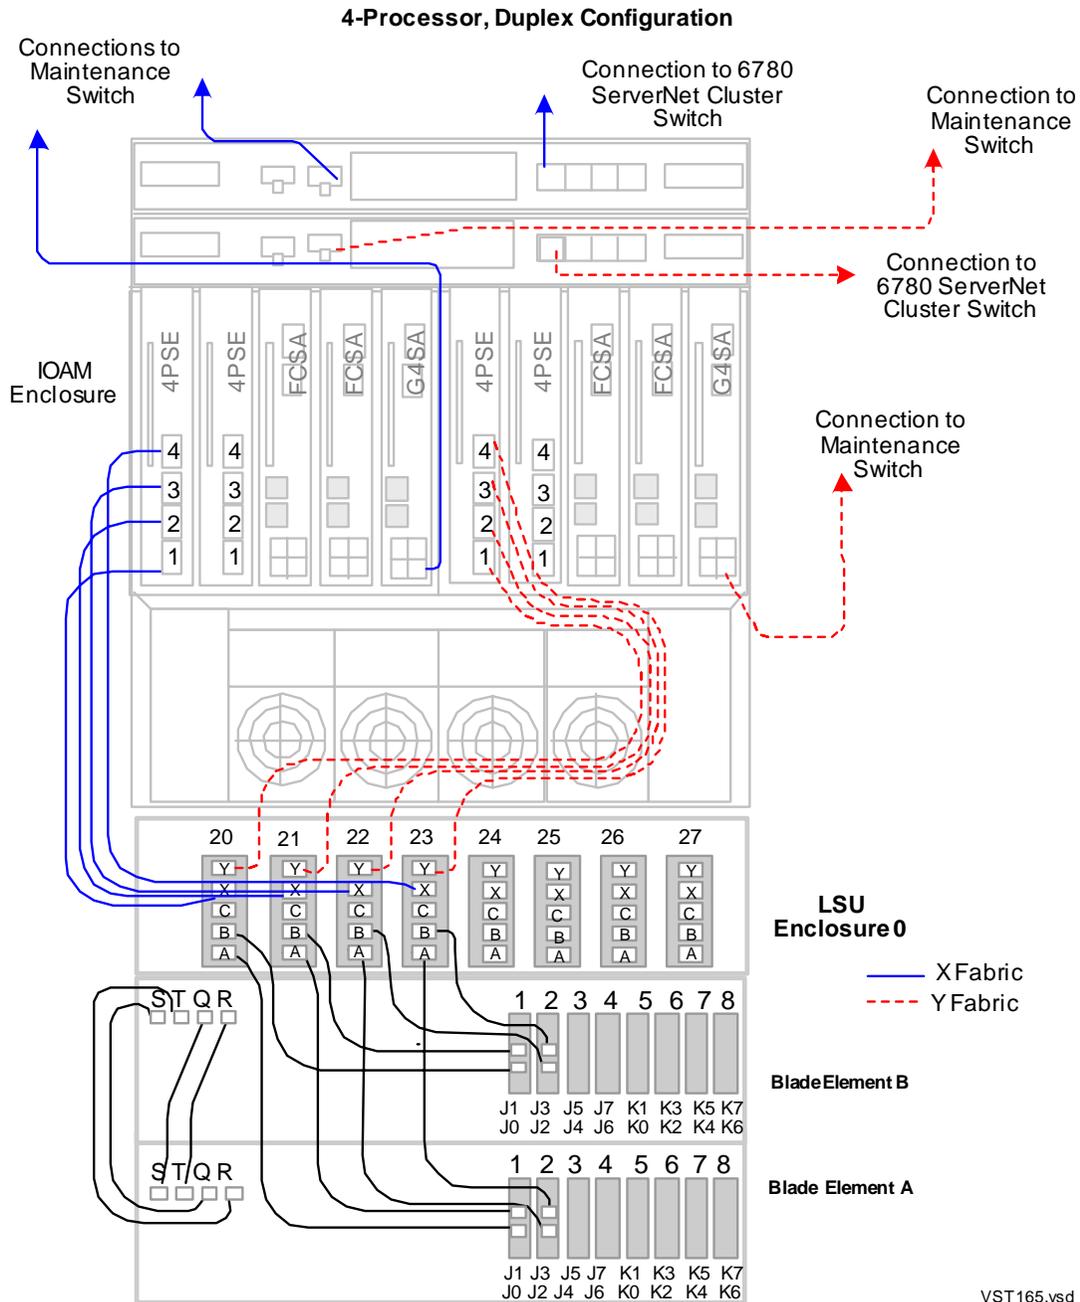
Figure 7-1. Integrity NonStop NS16000 System



Integrity NonStop NS14000 ServerNet Connectivity

ServerNet connections between I/O devices and processors depend on whether the Integrity NonStop NS14000 system has an IOAM enclosure or VIO enclosures. [Figure 7-2](#) shows an NS14000 system with an IOAM enclosure. For more information on Integrity NonStop NS14000 systems with VIO enclosures, see [Integrity NonStop NS14000 Systems](#) on page 2-3, the *NonStop NS14000 Planning Guide*, or the *Versatile I/O (VIO) Manual*.

Figure 7-2. Integrity NonStop NS14000 System with IOAM Enclosure



Integrity NonStop NS1000 ServerNet Connectivity

ServerNet connections between I/O devices and processors depend on whether the Integrity NonStop NS1000 system has an IOAM enclosure or VIO enclosures. For more information on Integrity NonStop NS1000 systems, see the *NonStop NS1000 Planning Guide*, *NonStop NS1000 Hardware Installation Manual*, or the *Versatile I/O (VIO) Manual*.

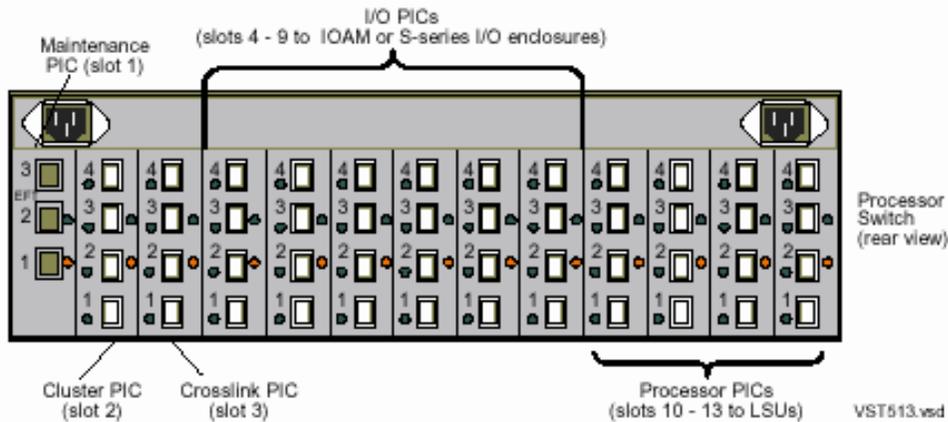
System I/O ServerNet Connections

For Integrity NonStop NS16000 systems, ServerNet connections to the system I/O devices (storage disk and tape drive as well as Ethernet communication to networks) radiate out from the p-switches for both the X and Y ServerNet fabrics.

ServerNet cables connected to the p-switch PICs in slots 10 through 13 come from the LSUs and processors. Cables connected to the PICs in slots 4 through 9 connect to one or more IOAM enclosures or to NonStop S-series I/O enclosures equipped with IOMF2 CRUs. [Figure 7-3](#) shows the connections to the PICs in a fully populated p-switch.

For Integrity NonStop NS14000 systems, see [Integrity NonStop NS14000 ServerNet Connectivity](#) on page 7-3. Like NS14000 systems, Integrity NonStop NS1000 systems use 4PSEs to provide ServerNet connections between I/O devices and processors. However, there are no LSUs; the 4PSEs connect directly to the Blade Elements. For more information, see the *NonStop NS1000 Hardware Installation Manual*.

Figure 7-3. I/O Connections to the PICs in a P-Switch



Monitoring the Status of the ServerNet Fabrics

The ServerNet fabrics provide the communication paths used for interprocessor messages, for communication between processors and I/O devices, and (in the case of

ServerNet clusters) for communication between systems. The ServerNet fabrics consist of two entirely separate communication paths—the X fabric and the Y fabric.

Note. If the system is a member of a ServerNet cluster, ServerNet connections to other members are accomplished by extending the ServerNet fabrics outside the system. Such external connections make up the external ServerNet fabrics. The *ServerNet Cluster Manual* provides additional information about monitoring the external ServerNet fabrics.

To monitor the status of the ServerNet fabrics:

- Use the OSM Service Connection to check the communication between processor enclosures, I/O enclosures, and systems.
- Use the Subsystem Control Facility (SCF) to check the status of interprocessor communication on the X and Y fabrics.

Monitoring the ServerNet Fabrics Using OSM

To check the ServerNet fabrics:

1. Log on to the OSM Service Connection.
2. Expand the tree pane to locate and select the internal or external ServerNet fabric objects:
 - a. The X and Y internal ServerNet fabric objects are located under the System and Fabric Group objects.
 - b. The X and Y external ServerNet fabric objects are located under the ServerNet Cluster object (if your system is part of a ServerNet cluster).
3. Check the fabric objects for:
 - a. If a fabric object icon is covered by a red or yellow triangular symbol, check the Attributes tab in the details pane for degraded attribute values. The Service State attribute is only displayed in the Attributes tab if it has a value of other than OK. If a degraded Service State is indicated, there will be an associated alarm to provide more information about the cause of the problem.
 - b. If a bell-shaped symbol is displayed next to a fabric icon in the tree pane, select the Alarms tab from the details pane. To get more information on an individual alarm, click to select the alarm, then right-click and select Details.
 - c. If a fabric object icon is covered by a yellow arrow, there is a problem on a subcomponent of that fabric. Expand the fabric object to locate the subcomponent object reporting the alarm or problem attribute.

For Integrity NonStop NS16000 systems, under each internal fabric object, you'll find its associated processor switch (p-switch) module and subcomponents. For Integrity NonStop NS14000 and NS1000 systems, there are no p-switches. Instead, ServerNet connectivity is provided through 4PSEs located in the IOAMs.

For more information on monitoring and recovery for external fabrics, see the appropriate ServerNet cluster manual for your particular ServerNet cluster configuration and hardware.

Monitoring the ServerNet Fabrics Using SCF

The SCF STATUS SERVERNET command displays a matrix for the ServerNet X fabric and a matrix for the ServerNet Y fabric. Each matrix shows the status of the paths between all pairs of processors.

Use the SCF STATUS SERVERNET command to display current information about the ServerNet fabric. At a TACL prompt:

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

```

l-> 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  UNA  UNA
01      UP  UP  UNA  UNA
02    <- DOWN
03    <- DOWN
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  UNA  UNA
01      UP  UP  UNA  UNA
02    <- DOWN
03    <- DOWN
04    <- DOWN
05    <- DOWN
06    <- DOWN
07    <- DOWN
08    <- DOWN
09    <- DOWN
10    <- DOWN
11    <- DOWN
12    <- DOWN
13    <- DOWN
14    <- DOWN
15    <- DOWN

```

In the preceding example of a 2-processor system:

- All ServerNet connections between processors 0 and 1 are up.
- Processors 2 through 15 do not exist on this system. As a result:
 - The status from processors 0 and 1 to processors 2 through 15 is displayed as unavailable (UNA) in both fabrics.
 - The status from processors 2 through 15 is displayed as down.

Normal ServerNet Fabric States

Normal states for a path on the ServerNet fabrics can be one of:

- UP
The path from the processor in the FROM row to the processor in the TO column is up. The status for all ServerNet connections between existing processors in a system should be UP.
- <-DOWN (for an entire row)
The processor in the FROM row is down or nonexistent. If the processor in the FROM row does not exist on your system, this status is normal. Otherwise, refer to [Identifying ServerNet Fabric Problems](#) on page 7-7.
- UNA (unavailable)
The processor in the TO column is down or nonexistent. Therefore, the path from the processor in the FROM row to the processor in the TO column is down. If the processor in the TO column does not exist on your system, this status is normal. Otherwise, refer to [Identifying ServerNet Fabric Problems](#) on page 7-7.

Identifying ServerNet Fabric Problems

Depending on how your system is configured, these states for a path on the ServerNet fabrics might indicate a problem:

- DIS (disabled)
The ServerNet fabric is down at the TO location. As a result, the path from the processor in the FROM row to the processor in the TO column is down for receiving; that is, the processor in the TO column cannot receive from any other processor or from I/O devices. DIS overrides both UP and DN.
- DN (down)
The path from the processor in the FROM row to the processor in the TO column is down because the path is failing. The processor in the FROM row cannot send to the processor in the TO column.
- <- DOWN (for an entire row)
The processor in the FROM row is down or nonexistent. For a processor that does exist on your system, this status is abnormal.
- ERROR *nnn* (for an entire row)
The processor in the FROM row unexpectedly returned a file-system error to that ServerNet fabric.
- UNA (unavailable)
The path from the processor in the FROM row to the processor in the TO column is down because the processor in the TO column is down. For a processor that does exist on your system, this status is abnormal. UNA overrides all other states.

Recovery Operations for the ServerNet Fabrics

For most recovery operations, refer to the *SCF Reference Manual for the Kernel Subsystem*.

Recovery Operations for a Down Disk Due to a Fabric Failure

When a path to a disk drive goes down due to a ServerNet fabric failure (either the ServerNet X or Y fabric is down), the storage subsystem automatically switches the paths of the disk drive if possible, so that the disk drive remains operational. This switching might result in a disk drive being placed in the STOPPED state with a substate of HARDDOWN.

You must restart any disk path that was using the fabric that went down. Otherwise the storage subsystem never attempts to use that path, which creates a potential single point of failure. For more information, refer to [Recovery Operations for a Down Disk or Down Disk Path](#) on page 10-14.

Recovery Operations for a Down Path Between Processors

When the status is either DIS (disabled) or DN (down), you can restart all paths between processors on the X fabric or Y fabric:

```
> SCF START SERVERNET $ZSNET.{X|Y}.*
```

Refer to the *SCF Reference Manual for the Kernel Subsystem*.

Recovery Operations for a Down Processor

If the status for an existing processor is <- DOWN or UNA, refer to [Recovery Operations for a Processor Halt](#) on page 9-9 for more information about recovery operations.

Recovery Operations for a File-System Error

For information about file-system errors, refer to the *Guardian Procedure Errors and Messages Manual*.

Related Reading

For more information about the interconnections between Integrity NonStop systems and NonStop S-series systems, see the *Integrity NonStop NS-Series Planning Guide*.

For more information about the ServerNet fabrics, see the *SCF Reference Manual for the Kernel Subsystem*.

I/O Adapters and Modules: Monitoring and Recovery

[When to Use This Section](#) on page 8-1

[I/O Adapters and Modules](#) on page 8-2

[Fibre Channel ServerNet Adapter \(FCSA\)](#) on page 8-2

[Gigabit Ethernet 4-Port Adapter \(G4SA\)](#) on page 8-2

[4-Port ServerNet Extender \(4PSE\)](#) on page 8-3

[Monitoring I/O Adapters and Modules](#) on page 8-3

[Monitoring the FCSAs](#) on page 8-4

[Monitoring the G4SAs](#) on page 8-5

[Monitoring the 4PSEs](#) on page 8-7

[Recovery Operations for I/O Adapters and Modules](#) on page 8-7

[Related Reading](#) on page 8-8

When to Use This Section

Use this section for monitoring and recovery information for the Fibre Channel ServerNet adapters (FCSAs) and the Gigabit Ethernet 4-port adapter. Information on ServerNet/DAs, the IOMF2 enclosure, and the I/O adapter module (IOAM) is available in NonStop S-series documentation.

For information about the disk drives or tape drives supported on a ServerNet/DA for your H-series RVU, refer to the *H-Series Highlights and Migration Planning Guide*.

Note. Starting with H06.08, new Integrity NonStop NS14000 and NS1000 servers are shipped with VIO enclosures instead of an IOAM enclosure. VIO enclosures provide the same functionality as IOAM enclosures. The monitoring and recovery principles described in this chapter are essentially the same for VIO enclosures; however, the components and OSM object names vary. For more information on VIO enclosures, see:

- *Versatile I/O (VIO) Manual*
 - *NonStop NS14000 (or NS1000) Planning Guide*
 - *NonStop NS14000 (or NS1000) Hardware Installation Manual*
 - *OSM Service Connection User's Guide*
-

I/O Adapters and Modules

Beginning with Integrity NonStop systems, interprocessor communications and I/O use dual ServerNet fabrics as a common interconnect means. Input/output components usually connect to the ServerNet fabrics through ServerNet adapters that are in an I/O adapter module (IOAM) enclosure. These adapters provide the system I/O to Fibre Channel storage devices and gigabyte Ethernet communications networks. Connections to the ServerNet fabric through a NonStop S-series I/O enclosure equipped with IOMF2s provide additional ServerNet interfacing for the Integrity NonStop I/O peripherals.

Even though the hardware architecture differs from one series of NonStop servers to another, the Integrity NonStop system can be networked with other NonStop systems using the same message system and the same network software.

Fibre Channel ServerNet Adapter (FCSA)

The FCSA provides Fibre Channel connectivity to certain external devices such as disk drives contained in a disk drive enclosure that supports fibre channel disks and an Enterprise Storage System (ESS).

Any connection between an Integrity NonStop system and a disk drive enclosure containing M8xxx fibre channel disks requires the services of two processes: the Fibre Channel Storage (FCS) Manager, which is part of the Storage Manager (\$ZZSTO), and the FCS Monitor (\$FCSMON), a persistent generic process that runs in all processors. An FCS Monitor process must be running in all processors. Each of the two SACs on an FCSA can support as many as four disk drive enclosures, for a total of eight per FCSA. The FCS Manager process assigns a SAC on an FCSA to a particular instance of the FCS Monitor.

Up to 10 FCSAs can be housed in an I/O adapter module (IOAM), which is mounted in an IOAM enclosure (except in Integrity NonStop NS14000 and NS1000 systems, where slot 1 is reserved for a [4-Port ServerNet Extender \(4PSE\)](#)). The form factor and connection technology of IOAM enclosures differ from the standard I/O enclosures that provide direct ServerNet access to external I/O devices. A pair of ServerNet switch boards, also located in the IOAM enclosure, provide connectivity between the processors and the FCSAs. All IOAM hardware can be monitored by OSM.

For information about the disk drives or tape drives supported through FCSAs for your H-series RVU, refer to the *H06.nn Release Version Update*.

Gigabit Ethernet 4-Port Adapter (G4SA)

The M8800 Gigabit Ethernet 4-port ServerNet adapter (G4SA) provides Gigabit connectivity between Integrity NonStop systems and Ethernet LANs. G4SAs are installed in slots 1 through 5 of an I/O adapter module (IOAM) (except in Integrity NonStop NS14000 and NS1000 systems, where slot 1 is reserved for a [4-Port ServerNet Extender \(4PSE\)](#)). There are two IOAMs in an IOAM enclosure, so a total of 10 G4SAs can be installed in an 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 in a NonStop S-series enclosure.

A G4SA has three primary system connections:

- Data transfer interface (ServerNet)
- Maintenance entity (ME) interface
- Power interface

The data transfer interface consists of ports to the ServerNet X and Y fabrics. The ports connect to the ServerNet addressable controller (SAC) on the adapter. If one ServerNet fabric fails, the G4SA can still be accessed using the remaining fabric.

The Maintenance Entity (ME) interface contains the circuitry required to meet the maintenance system requirements of an active-logic adapter.

The G4SA receives power through a shielded, high-density, metric connector module. The connector module provides attachments to the two ServerNet fabrics.

G4SAs are configured and managed through the Subsystem Control Facility (SCF) interface to the ServerNet LAN Systems Access (SLSA) subsystem. The SLSA subsystem is preinstalled and preconfigured and is started during the system-load sequence.

For information about the SLSA subsystem, refer to the *LAN Configuration and Management Manual*.

4-Port ServerNet Extender (4PSE)

A component in Integrity NonStop NS14000 and NS1000 systems only, 4PSEs provide ServerNet connectivity between processors and the IOAM enclosure (functionality provided by p-switches in an Integrity NonStop NS16000 system). 4PSEs are located in slot one (and optionally slot 2) of each IOAM. They are connected to the processors through LSUs in Integrity NonStop NS14000 systems, directly to the processors (with no LSUs) in Integrity NonStop NS1000 systems. FCSAs and G4SAs can be installed in slots 2 through 5 of the two IOAMs in the IOAM enclosure for communications to storage devices and subsystems as well as to LANs.

See [Monitoring the 4PSEs](#) on page 8-7.

Monitoring I/O Adapters and Modules

Use the Subsystem Control Facility (SCF) or the OSM Service Connection to monitor the FCSAs, G4SAs, and ServerNet/DAs.

Monitoring the FCSAs

For a general, top-down approach for using OSM to monitor system components, refer to [Using OSM to Monitor the System](#) on page 3-7.

To monitor the FCSA and its attached devices with SCF, use the SCF INFO and SCF STATUS commands.

For example, to monitor all FCSAs using SCF:

```
> SCF STATUS ADAPTER $ZZSTO.#FCSA*, DETAIL
```

The *SCF Reference Manual for the Storage Subsystem* provides reference details and examples for using the SCF INFO and SCF STATUS commands.

When monitoring FCSAs using the OSM Service Connection, the states of the FCSAs should indicate normal operation. [Table 8-1](#) lists the possible states for the FCSA.

Table 8-1. Service, Flash Firmware, Flash Boot Firmware, Device, and Enabled States for the FCSA (page 1 of 2)

State	Description
Service State: OK	The resource is functioning normally and does not require attention or service.
Service State: Attention Required	The resource requires operator attention. This condition sometimes generates an alarm, and the component that requires attention is colored yellow in the tree pane and in the Physical and ServerNet views of the view pane.
Service State: Service Required	The resource requires service. This condition generates an alarm, and the component that requires service is colored red in the tree pane and in the Physical and ServerNet views of the view pane.
Flash Compare State: Same	Current and default versions are the same.
Flash Compare State: Up-rev	The current version is newer than the default version.
Flash Compare State: Down-rev	The default version is newer than the current version.
Flash Compare State: Unknown	Unable to compare files.
Device State: Aborting	Processing is terminating.
Device State: Defined	State is defined by the NonStop OS.
Device State: Degraded	Performance is degraded.
Device State: Diagnose	A diagnostic test is running on the component.

Table 8-1. Service, Flash Firmware, Flash Boot Firmware, Device, and Enabled States for the FCSA (page 2 of 2)

State	Description
Device State: Initializing	Processing is starting up.
Device State: Not Configured	The component is not configured.
Device State: Started	The component is running.
Device State: Starting	Processing is starting up.
Device State: Stopped	Processing has been terminated.
Device State: Stopping	Processing is being terminated.
Device State: Unknown	Component is not responding.
Device State: OK	Component is accessible.
Enabled State: Disabled	The component is present but not operational, possibly because the Disable action was performed.
Enabled State: Enabled	The component is operational.
Enabled State: Fault	A problem was detected. The component might be functioning below standard or not responding. Inspect the component.
Enabled State: Off	The component is not functional.
Enabled State: On	The component is functional.
Enabled State: Unknown	State is unknown; component might not be responding.

Monitoring the G4SAs

Use the Subsystem Control Facility (SCF) or the OSM Service Connection to monitor the G4SAs.

For a general, top-down approach for using OSM to monitor system components, refer to [Using OSM to Monitor the System](#) on page 3-7.

To monitor the G4SA and its attached devices with SCF, use the SCF INFO and SCF STATUS commands.

For example, to monitor G4SAs using SCF:

```
> SCF STATUS ADAPTER $ZZLAN.G1123
```

The *LAN Configuration and Management Manual* provides reference details and examples for using the SCF INFO and SCF STATUS commands.

When monitoring G4SAs using the OSM Service Connection, the states of the G4SAs should indicate normal operation. [Table 8-2](#) lists the possible states for the G4SA.

Table 8-2. Service, Device, and Enabled States for the G4SA (page 1 of 2)

State	Description
Service State: OK	The resource is functioning normally and does not require attention or service.
Service State: Attention Required	The resource requires operator attention. This condition sometimes generates an alarm, and the component that requires attention is colored yellow in the tree pane and in the Physical and ServerNet views of the view pane.
Service State: Service Required	The resource requires service. This condition generates an alarm, and the component that requires service is colored red in the tree pane and in the Physical and ServerNet views of the view pane
Device State: Aborting	Processing is terminating.
Device State: Defined	State is defined by the NonStop OS.
Device State: Degraded	Performance is degraded.
Device State: Diagnose	A diagnostic test is running on the component.
Device State: Initializing	Processing is starting up.
Device State: Not Configured	The component is not configured.
Device State: Started	The component is running.
Device State: Starting	Processing is starting up.
Device State: Stopped	Processing has been terminated.
Device State: Stopping	Processing is being terminated.
Device State: Unknown	Component is not responding.
Device State: OK	Component is accessible.
Enabled State: Disabled	The component is present but not operational, possibly because the Disable action was performed.
Enabled State: Enabled	The component is operational.
Enabled State: Fault	A problem was detected. The component might be functioning below standard or not responding. Inspect the component.

Table 8-2. Service, Device, and Enabled States for the G4SA (page 2 of 2)

State	Description
Enabled State: Off	The component is not functional.
Enabled State: On	The component is functional.
Enabled State: Unknown	State is unknown; component might not be responding.

Monitoring the 4PSEs

For a general, top-down approach for using OSM to monitor system components, refer to [Using OSM to Monitor the System](#) on page 3-7.

For 4PSEs, the OSM Service Connection reports, in the form of attributes, the

- Service State (displayed only if the value is something other than OK) – When a negative Service State value is displayed, check for alarms and see alarm details for the probable cause and suggested repair actions.
- Power State
- Device State

OSM actions allow you to power on or off a 4PSE, turn the LED on or off, and the Replace action launches a documented service procedure to guide you through replacement.

Recovery Operations for I/O Adapters and Modules

- Examine the contents of the event message log for the subsystem. For example, the ServerNet LAN Systems Access (SLSA) subsystem or Storage subsystem might have issued an event message that provides information about the resource failure. Event messages returned by the SLSA and Storage subsystems are described in the SLSA and STO sections of the Operator Messages Manual, respectively.
- HP provides a comprehensive library of troubleshooting guides for the communications subsystems. Attempt to analyze the problems and restart the process or object using the commands described in the appropriate manual listed in [Related Reading](#) on page 8-8. If you are unable to start a required process or object, contact your service provider.

Related Reading

For more information about monitoring and performing recovery operations for the I/O adapters and the SLSA and Storage subsystems, see the manuals listed in [Table 8-3](#). The appropriate manual to use depends on how your system is configured.

Table 8-3. Related Reading for I/O Adapters and Modules

For Information About	Refer to...
Monitoring and recovery information for the FCSA	<i>Fibre Channel ServerNet Adapter Installation and Support Guide</i>
	<i>SCF Reference Manual for the Storage Subsystem</i>
	<i>OSM Service Connection User's Guide</i>
	OSM online help
Monitoring and recovery information for the G4SA	<i>Gigabit Ethernet 4-Port Adapter Installation and Support Guide</i>
	<i>LAN Configuration and Management Manual</i>
	<i>OSM Service Connection User's Guide</i>
	OSM online help
Monitoring and recovery information for the ServerNet/DA	<i>6760 ServerNet/DA Manual</i>
	<i>SCF Reference Manual for the Storage Subsystem</i>
	<i>OSM Service Connection User's Guide</i>
	OSM online help
General information on the IOAM	<i>NonStop NS16000 Hardware Installation Manual</i>

Processors and Components: Monitoring and Recovery

[When to Use This Section](#) on page 9-1

[Overview of the NonStop Blade Complex](#) on page 9-2

[Monitoring and Maintaining Processors](#) on page 9-4

[Monitoring Processor Status Using the OSM Low-Level Link](#) on page 9-5

[Monitoring Processor Status Using the OSM Service Connection](#) on page 9-5

[Monitoring Processor Performance Using ViewSys](#) on page 9-7

[Identifying Processor Problems](#) on page 9-7

[Processor or System Hangs](#) on page 9-7

[Processor Halts](#) on page 9-8

[OSM Alarms and Attribute Values](#) on page 9-8

[Recovery Operations for Processors](#) on page 9-9

[Recovery Operations for a Processor Halt](#) on page 9-9

[Halting One or More Processors](#) on page 9-10

[Reloading a Single Processor on a Running Server](#) on page 9-10

[Recovery Operations for a System Hang](#) on page 9-14

[Enabling/Disabling Processor and System Freeze](#) on page 9-15

[Freezing the System and Freeze-Enabled Processors](#) on page 9-15

[Dumping a Processor to Disk](#) on page 9-15

[Backing Up a Processor Dump to Tape](#) on page 9-19

[Replacing Processor Memory](#) on page 9-19

[Replacing the Processor Board and Processor Entity](#) on page 9-19

[Submitting Information to Your Service Provider](#) on page 9-19

[Related Reading](#) on page 9-22

When to Use This Section

Use this section to monitor processors and to perform recovery operations such as processor dumps.

Overview of the NonStop Blade Complex

Note. This section does not apply to Integrity NonStop NS1000 systems, which use the NSVA rather than NSAA architecture (see [NonStop System Architectures](#) on page 2-2). For more information on Integrity NonStop NS1000 systems, see [Differences Between Integrity NonStop NS-Series Systems](#) on page 2-2, the *NonStop NS1000 Planning Guide*, or the *NonStop NS1000 Hardware Installation Manual*.

The basic building block of the modular NonStop advanced architecture (NSAA) compute engine is the NonStop Blade Complex, which consists of two or three processor modules called NonStop Blade Elements. Each Blade Element houses two or four microprocessors called processor elements (PEs). A logical processor consists of one processor element from each Blade Element. Although a logical processor physically consists of multiple processor elements, it is convenient to think of a logical processor as a single entity within the system. Each logical processor has its own memory, its own copy of the operating system, and processes a single instruction stream. NSAA logical processors are usually referred to simply as “processors.”

All input and output to and from each NonStop Blade Element goes through a logical synchronization unit (LSU). The LSU interfaces with the ServerNet fabrics and contains logic that compares all output operations of a logical processor, ensuring that all NonStop Blade Elements agree on the result before the data is passed to the ServerNet fabrics.

A processor with two NonStop Blade Elements comprise the dual modular redundant (DMR) NonStop Blade Complex, which is also referred to as a duplex system. This duplex system provides data integrity and system availability that is comparable to NonStop S-series systems, but at considerably faster processing speeds.

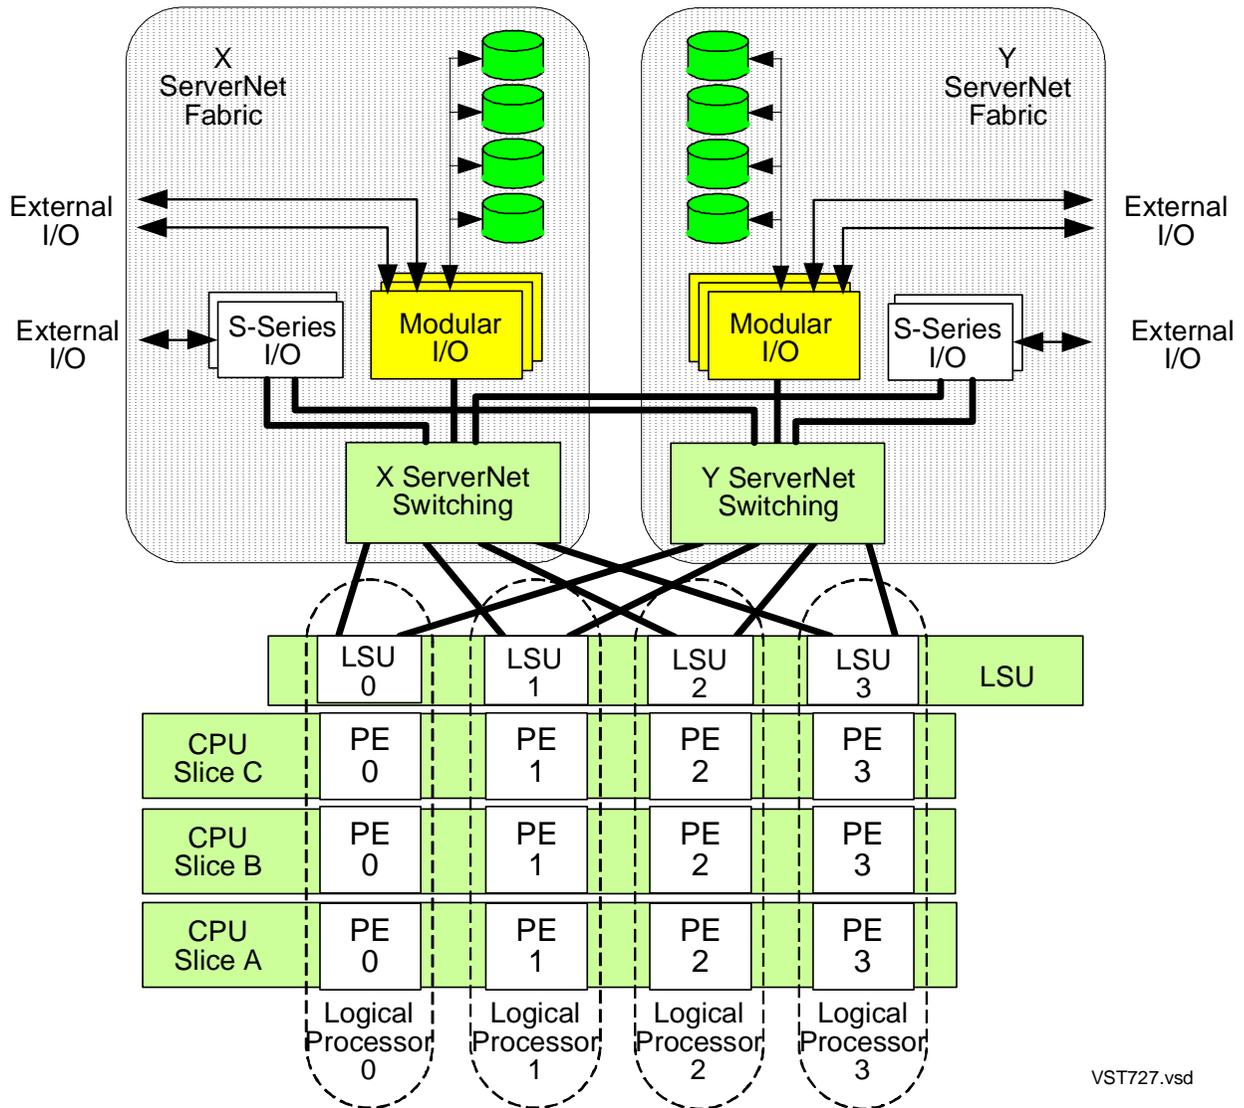
Three NonStop Blade Elements plus their associated LSUs make up the triple modular redundant (TMR) NonStop Blade Complex, which is referred to as a triplex system. The triplex system provides the same processing speeds as the duplex system while also enabling hardware fault recovery that is transparent to all but the lowest level of the NonStop operating system (OS).

In the event of a processor fault in either a duplex or triplex system, the failed component within a NonStop Blade Element (processor element, power supply, and so forth) can be replaced while the system continues to run. A single Integrity NonStop system can have up to four NonStop Blade Complexes for a total of 16 processors. Processors communicate with each other and with the system I/O over dual ServerNet fabrics.

A ServerNet fabric is a complex web of links that provide a large number of possible paths from one point to another. Two communications fabrics, the X and Y ServerNet fabrics, provide redundant, fault-tolerant communications pathways. If a hardware fault occurs on one of the ServerNet fabrics, communications continues on the other with hardware fault recovery transparent to all but the lowest level of the OS.

[Figure 9-1](#) is an overview of the modular NSAA and shows one NonStop Blade Complex with four processors, the I/O hardware and the ServerNet fabrics.

Figure 9-1. Modular NSAA With One NonStop Blade Complex and Four Processors



For Integrity NonStop NS16000 servers, ServerNet communications are controlled by processor switch (or p-switch) modules; one for each of the ServerNet fabrics (represented by X and Y ServerNet Switching blocks in [Figure 9-1](#)). P-switches connect input/output components to the NS-series processors. I/O components usually tap into the ServerNet fabrics through ServerNet adapters in IOAMs (represented by Modular I/O blocks in the diagram). These adapters provide the system I/O to fibre channel storage devices and gigabyte Ethernet communications networks. However, P-switches also provide connectivity for legacy I/O through I/O enclosures equipped with IOMF2s (represented by S-Series I/O blocks in the diagram).

Integrity NonStop NS14000 systems do not have p-switches and cannot be connected to legacy NonStop S-series I/O enclosures. For more information, see [Integrity NonStop NS14000 ServerNet Connectivity](#) on page 7-3).

In summary, these terms describe the NSAA processor:

Term	Description
Blade Complex	Consists of two Blade Elements (in a duplex system) or three Blade Elements (in a triplex system) and up to four logical processors and their associated LSUs. An Integrity NonStop system includes up to four Blade Complexes.
Blade Element	Consists of a chassis, processor board containing two or four PEs (one representing each logical processor in the Blade Complex), memory, I/O interface board, midplane, optics adapters, fans, and power supplies. Blade Elements are mounted in a 19-inch computer equipment rack.
Processor element (PE)	A single Itanium microprocessor with its associated memory. A PE is capable of executing an individual instruction stream and I/O communication through fiber-optic links.
Logical processor	One PE from each Blade Element executing a single instruction stream. A duplex system has two PEs forming a logical processor. A triplex system has three PEs.

Monitoring and Maintaining Processors

To monitor processors, use OSM, the ViewSys product, and other tools. Monitoring and maintaining processors includes:

- [Monitoring Processors Automatically Using TFDS](#) on page 9-4
- [Monitoring Processor Status Using the OSM Low-Level Link](#) on page 9-5
- [Monitoring Processor Status Using the OSM Service Connection](#) on page 9-5
- [Monitoring Processor Performance Using ViewSys](#) on page 9-7
- [Monitoring EMS Event Messages](#) on page 4-1

Monitoring Processors Automatically Using TFDS

HP Tandem Failure Data System (TFDS) should be used to proactively monitor processors and manage processor halts. Configured and running before a halt occurs, TFDS can help determine the type of recovery operation needed and:

- If TFDS determines that the entire processor should be dumped be reloading, it automatically dumps, then reloads the processor.
- If TFDS determines that only the processor element (PE) for one Blade Element needs to be dumped, it reloads the processor, excluding that Blade Element, dumps the Blade Element, then reintegrates it back into the running processor.
- Collects the files necessary to analyze the problem.
- Sends halt information message to the EMS collector. If configured in OSM, a dial-out message is sent to HP Global Support to notify them of the halt.

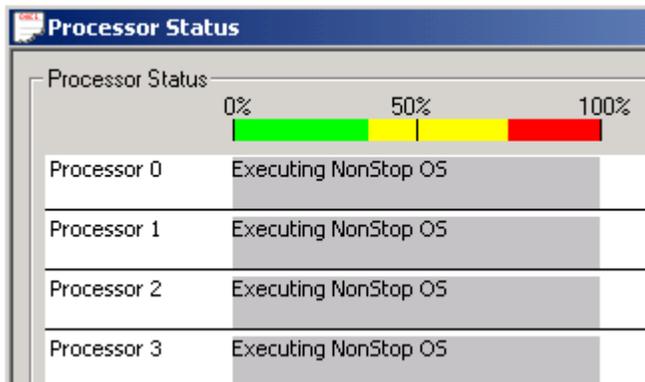
For more information on configuring and using TFDS, see the *Tandem Failure Data System (TFDS) Manual*.

Monitoring Processor Status Using the OSM Low-Level Link

From the OSM Low-Level Link, use the Processor Status dialog box to determine if the processors are running:

1. Log on to the OSM Low-Level Link.
2. On the toolbar, click the **Processor Status** button.
3. The status for all processors should be “Executing NonStop OS.” (See [Figure 9-2.](#)) If not, refer to [Identifying Processor Problems](#) on page 9-7.

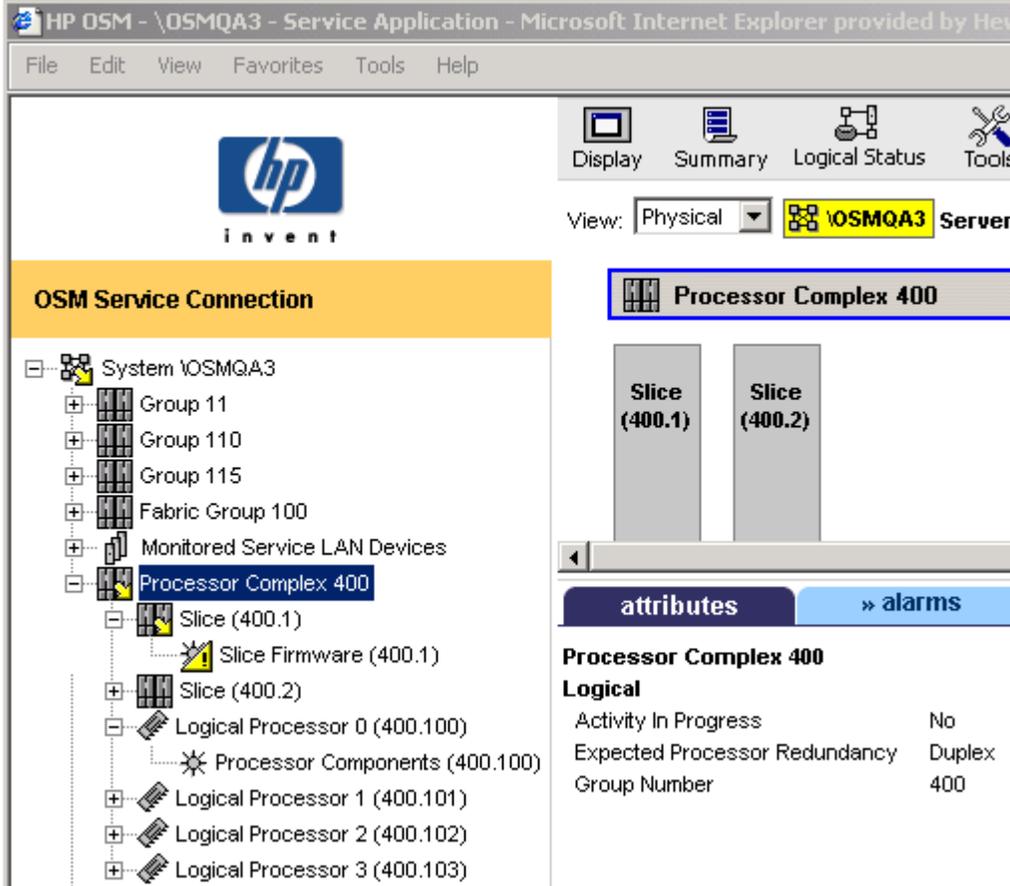
Figure 9-2. Processor Status Display



Monitoring Processor Status Using the OSM Service Connection

For Integrity NonStop NS-series systems, the OSM Service Connection displays processor-related components under Processor Complex objects in the tree pane. There can be up to four Processor Complex objects per NS-series system, for a total of 16 processors. When you expand a Processor Complex object (see [Figure 9-3](#)), you should see two or three Blade Element objects and either two or four Logical Processor objects.

Under the each Blade Element object, a Blade Element Firmware object displays the firmware version information for that Blade Element. Under each Logical Processor object, a Processor Components object represents (provides attributes and actions for) the processor's associated logical synchronization unit (LSU) and processor elements (PEs). From the processor perspective, each PE is identified as A, B, or C (to identify the Blade Element it is associated with); whereas, from the Blade Element perspective, each PE is identified as 0, 1, 2, or 3 (to identify the Logical Processor it is associated with).

Figure 9-3. OSM Representation of Processor Complex

To check processor-related components using the OSM Service Connection:

1. Expand the tree pane to check all Processor Complex objects.
2. If a Processor Complex object icon contains a yellow arrow (as illustrated in [Figure 9-3](#)), expand that complex to check its subcomponents.
 - If any processor subcomponent is displaying a red or yellow triangular symbol over its object icon, check the Attributes tab for degraded attribute values.
 - If a bell-shaped alarm icon appears next to the subcomponent's object icon, check the Alarms tab. To get details on an alarm, select, then right-click on the alarm and select **Details**.

If a problem exists on a logical processor, the Halt Flag attribute has a value of true, and a Halt Code attribute value is displayed, refer to the *Processor Halt Codes Manual*.

Monitoring Processor Performance Using ViewSys

Use the ViewSys product to view system resources online and to see information on system performance. ViewSys provides information about processor activity. Using ViewSys, you can list the processors on your system and determine their status. For more information, refer to [ViewSys on page B-6](#).

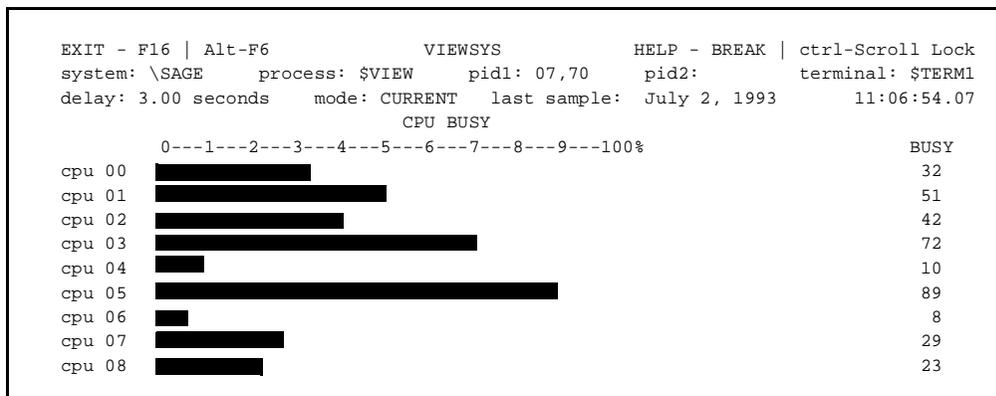
To use ViewSys to obtain information about processor activity, at a TACL prompt:

```
> VIEWSYS
```

A series of bar graphs that summarize processor performance statistics appears on your terminal.

Note. The Measure utility also collects and displays statistics about system performance and the performance of processors and other system components. Operations management personnel often use this utility to help fine-tune and balance a system. For instructions on using this utility, refer to the *Measure User's Guide* and the *Measure Reference Manual*.

After the first ViewSys screen appears, press F1 to view processor busy statistics:



To exit ViewSys, press F16.

Identifying Processor Problems

Processor problems include system hangs, processor halts, OSM alarms.

Processor or System Hangs

A processor hang occurs when system components wait for an event that is not going to happen. An unexpected event such as a deadlock (two or more processors waiting for each other) might have occurred. In some cases, the entire system might hang, or be unresponsive.

Processor Halts

When certain errors occur (such as when data integrity is at risk), the operating system cannot correct the problem and must halt all application and system processes running in the associated processor. The remaining running processors in the system each send a message, reporting the halted processor as down. The other processors in the system, including the backup to the halted processor, are not affected by the errors that caused the processor to halt unless they are freeze-enabled. Two types of processor halts display a processor halt code in the Processor Status dialog box:

- A halt instruction results in a processor halt.

When the operating system detects a millicode or software error that it cannot correct, it can execute a halt instruction to suspend all application and system processes running in the associated processor. The status of the halted processor becomes:

```
Halt code = %nnnnnn
```

Unlike a freeze instruction, a halt instruction affects only one processor.

- A processor can be halted by a freeze instruction.

A freeze-enabled processor can be frozen by another frozen processor. When a freeze instruction is executed, any processors that are freeze-enabled also freeze immediately. When the operating system detects a software error that it cannot correct, it can execute a freeze instruction to suspend all application and system processes running in the associated processor. The status of the frozen processor becomes:

```
Freeze code = %nnnnnn
```

If system freeze is enabled, the status for all other freeze-enabled processors becomes:

```
Frozen by other processor
```

The *Processor Halt Codes Manual* documents processor halt codes.

Note. Do not freeze-enable a processor unless instructed to do so by your service provider.

OSM Alarms and Attribute Values

When the OSM Service Connection displays processor-related alarms or problem attributes, the alarm and attribute details often indicate the appropriate recovery action. Recovery actions might call for:

- Trying the [Recovery Operations for Processors](#) on page 9-9.
- Contacting your service provider.

Recovery Operations for Processors

The architecture of Integrity NonStop NS-series servers offers recovery options not available in NonStop S-series. Because a logical processor has a physical processor element (PE) in each of up to three NonStop Blade Elements, in some cases, you no longer have to choose between taking the time to dump entire processors and skipping the dump to reload the system as quickly as possible. In those cases, you can reload a halted processor immediately, while excluding the PE for one Blade Element, then dump that excluded PE before reintegrating the Blade Element into the running processor.

Note. The parts of this section that do not apply to Integrity NonStop NS1000 systems include all references to processor elements (PEs) the RELOAD command OMITBLADE option. For more information on Integrity NonStop NS1000 systems, see [Differences Between Integrity NonStop NS-Series Systems](#) on page 2-2, the *NonStop NS1000 Planning Guide*, or the *NonStop NS1000 Hardware Installation Manual*.

Processor recovery operations for your NS-series system might include:

- [Recovery Operations for a Processor Halt](#) on page 9-9
- [Halting One or More Processors](#) on page 9-10
- [Reloading a Single Processor on a Running Server](#) on page 9-10
- [Recovery Operations for a System Hang](#) on page 9-14
- [Enabling/Disabling Processor and System Freeze](#) on page 9-15
- [Freezing the System and Freeze-Enabled Processors](#) on page 9-15
- [Dumping a Processor to Disk](#) on page 9-15
- [Backing Up a Processor Dump to Tape](#) on page 9-19
- [Replacing Processor Memory](#) on page 9-19
- [Replacing the Processor Board and Processor Entity](#) on page 9-19
- [Submitting Information to Your Service Provider](#) on page 9-19

Recovery Operations for a Processor Halt

HP Tandem Failure Data System (TFDS) should be used to proactively monitor processors and manage processor halts. Configured and running before a halt occurs, TFDS can help determine and perform the type of recovery operation needed (see [Monitoring Processors Automatically Using TFDS](#) on page 9-4). For information on configuring and using TFDS, see the *Tandem Failure Data System (TFDS) Manual*.

If all processors have halted (the system is down), TFDS cannot perform an automatic dump or reload. You must load the system, as described in [Performing a System Load From a Specific Processor](#) on page 15-11. You can omit one Blade Element from the load operation, to dump after the system is running. You can also dump the remaining processors as needed—dump the entire processor before reloading, or reload and omit Blade Element to dump later. For more information, see [Dumping a Processor to Disk](#) on page 9-15.

The best manual recovery operation for your processor halt scenario depends on the type of halt. You should record the halt information (as described in [Identifying Processor Problems](#) on page 9-7), and contact your service provider to help you determine the appropriate course of action.

This section describes the various options for [Dumping a Processor to Disk](#) and [Reloading a Single Processor on a Running Server](#) that you might be directed to use.

Halting One or More Processors

To place a selected processor or processors in a halt state and set the status and registers of the processor or processors to an initial state:

1. Log on to the OSM Low-Level Link.
2. On the toolbar, click **Processor Status** button.
3. In the Processor Status dialog box, select the processor to be halted or select all the processors to halt all of them.
4. Select **Processor Actions>Halt**.
5. Click **Perform action**.
6. A message box appears and asks whether you are sure you want to perform a halt on the selected processor or processors. Click **OK**.

Reloading a Single Processor on a Running Server

Sometimes one or more processors in a running server are not operating. For information on how to determine whether a processor is operating, see [Monitoring Processor Performance Using ViewSys](#) on page 9-7.

Unlike NonStop S-series, you don't always have to wait until you dump the processor before reloading it. This section describes how, when appropriate, to exclude the processor element (PE) from one NonStop Blade Element during the reload operation, so you can get the rest of the processor running, take the dump, then reintegrate the PE back into the running processor.

After you have determined that a processor is not operating, check that the processor is halted. If it needs to be halted, see [Halting One or More Processors](#) on page 9-10).

Collect information about the reason for the halt (as described in [Identifying Processor Problems](#) on page 9-7) to send to your service provider along with the dump file. In the Low-Level Link Processor Status dialog box, write down the halt code and status message for the processor.

Your options for reloading a processor on a running server are:

- [Using TACL RELOAD to Perform Reload on page 9-11](#)
- [Using the OSM Service Connection to Perform Reload on page 9-13](#)

Following the reload, if TFDS is not configured to take the dump automatically, you can perform a dump of the omitted PE while normal operations resume on the reloaded PEs within that logical processor. See the [Dumping a Processor to Disk](#) on page 9-15.

Using TACL RELOAD to Perform Reload

Run the RELOAD utility to reload the remaining processors after the first processor in a system has been brought up, or to recover a processor that has failed.

The H-series RELOAD utility allows you to omit (or exclude) a Blade Element from the reload operation. This allows you get the processor running for the PEs on the other Blade Elements, take a dump of the PE on the omitted Blade Element, and then reintegrate it back into the running processor. The OMITBLADE parameter allows you to specify the Blade Element (A, B, or C) to be excluded from the reload or, used without specifying a particular Blade Element, OMITBLADE selects a Blade Element to be excluded.

1. Select **File>Start Terminal Emulator**.
2. From the menu, select **For Startup TACL**. Two OutsideView windows appear.
3. In the Enter Telnet IP Address box, type the IP address that is currently configured for your primary service connection. If the processor for the primary service connection IP address is down, type the IP address that is currently configured for your backup service connection.
4. Click **OK**. The OutsideView window is active, and the TACL1> prompt appears.
5. Enter the super ID (255,255) and press Enter.
6. Type the password and press Enter.

A \$SYSTEM STARTUP 1> prompt appears (the prompt depends on your defaults).

7. Use the RELOAD command as appropriate for your scenario.

Note. If you plan to dump the PE for one Blade Element after reloading, use RELOAD with the OMITBLADE parameter. If you don't know which Blade Element to specify, use OMITBLADE without specifying A, B, or C. RELOAD will choose an appropriate Blade Element, and reply with the letter of the Blade Element that was omitted. Use this to specify the Blade Element in the RCVDUMP command.

```
RELOAD [ / run-option [ , run-option ] ... / ]
        cpu-set [ ; cpu-set ] ...
```

run-option

is any of the options described in the RUN[D|V] Command, in the *HP NonStop TACL Reference Manual*.

cpu-set

is a set of processors (and options) to be reloaded. Specify *cpu-set* as:

```
{ cpu-range } [ , option, option, ... ]
{ ( cpu-range, cpu-range, ... ) }
{ * }
```

cpu-range

is one of these:

cpu
cpu-cpu

cpu

is the processor number, an integer from 0 through 15.

cpu-cpu

is two processor numbers separated by a hyphen, specifying a range of processors. In a range specification, the first processor number must be less than the second.

option

is one of these:

NOSWITCH
[PRIME|NOPRIME]
fabric
OMITBLADE [A|B|C]
\$volume [.sysnn.osdir]

NOSWITCH

specifies that, when a processor is reloaded, there is no default autoswitch of controller ownership to the configured primary processor.

[PRIME|NOPRIME]

Sets up the logical processor for the reload operation. NOPRIME is the default.

fabric

specifies whether the X fabric or Y fabric is used for the transfer of the operating system image to the processor during the RELOAD operation.

0 = X fabric

1 = Y fabric

The default option is the X fabric.

OMITBLADE [A|B|C]

The PE on the selected Blade Element will not be reloaded when other PEs in that logical processor are reloaded. If you do not provide an argument (A or B or C) the system will choose a Blade Element to omit.

\$volume [.sysnn.osdir]

specifies a volume other than \$SYSTEM where the operating system image (*sysnn.osdir*) to be used for reloading the processor is located.

Depending on your requirements, this specification could take the form of *\$volume* or *\$volume.sysnn.osdir* or *sysnn.osdir* or *osdir*.

*

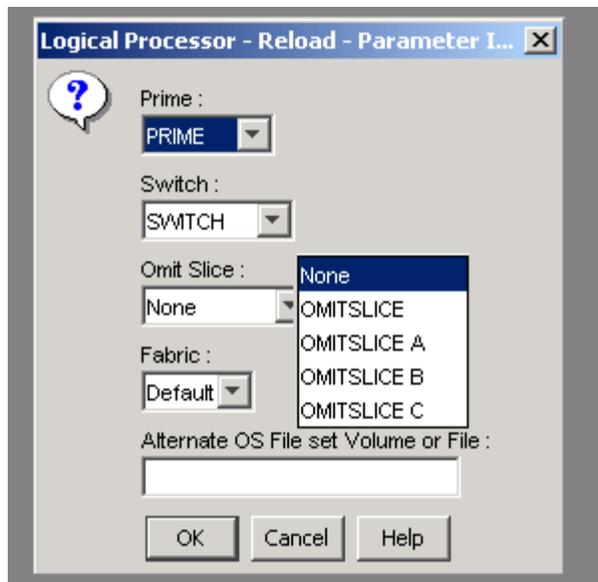
specifies that all failed processors should be reloaded.

8. Check the OutsideView window for status messages, which will report successes or errors during the load. Monitor the state of the processor you are loading until it is executing the NonStop Kernel operating system.
9. If the load fails, check the parameters and reload the processor. If the load fails again, contact your service provider.

Using the OSM Service Connection to Perform Reload

The OSM Service Connection provides a Reload action on the Logical Processor object. You can perform the action on a single or multiple processors. The OSM action lets you reload an entire processor or omit a Blade Element from the reload action so you can dump the PE for that Blade Element before reintegrating it into the running processor.

- To reload a single processor:
 1. Select the Logical Processor object for the processor you want to reload.
 2. Right-click and select **Actions**.
 3. Select **Reload**, click **Perform action**.
 4. Click **OK** to dismiss the confirmation dialog box.
 5. In the Logical Processor Reload Parameters dialog box, select the appropriate options. See OSM online help for information about the options.
 6. Click **OK**.



- To reload a multiple processors, use the Multi-Resource Actions dialog box (available from the Display menu of the OSM Service Connection).
 1. In the Multi-Resource Actions dialog box, select **Logical Processor** from the Resource Types list. All Logical Processors in the system will be displayed in the right-hand pane.
 2. Select **Reload** from the Action list under Selection Criteria.
 3. From the list of Logical Processors, select the ones you want to reload and click **Add** to move them to the lower list (you can select and add one at a time, or Ctrl-click to select more than one).
 4. Once all the processors you want to reload (and only those you want to reload) are in the lower list, click **Perform Action**.
 5. Click **OK** to the dismiss the confirmation dialog box.
 6. In the Logical Processor Reload Parameters dialog box (same dialog box pictured in procedure for reloading a single processor), select the appropriate options. See OSM online help for information about the options.
 7. Click **OK**.

Recovery Operations for a System Hang

If a system hang occurs on an NS-series server, you no longer have to choose between taking the time to dump the system or quickly loading the system without the benefit of getting the dump information. You can now load processor 0 or 1 while excluding one processor element (PE) in that processor. This allows you to reload the rest of the processors, get the system running as soon as possible, then take a dump of the PE that was excluded from the load operation to send to your service provider for analysis of the problem.

After a system hang has occurred, under the direction of your service provider, recovery operations might include:

1. Enabling System Freeze (see [Enabling/Disabling Processor and System Freeze](#) on page 9-15).
2. [Freezing the System and Freeze-Enabled Processors](#) on page 9-15.
3. Start the system by loading Processor 0 or 1, as described in [Performing a System Load From a Specific Processor](#) on page 15-11. You can omit one Blade Element from the load operation, to dump after the system is running. You can also dump the remaining processors as needed—dump the entire processor before reloading, or reload and omit Blade Element to dump later. For more information, see [Dumping a Processor to Disk](#) on page 9-15.
4. Disabling System Freeze (see [Enabling/Disabling Processor and System Freeze](#) on page 9-15).

Enabling/Disabling Processor and System Freeze

- △ **Caution.** Enabling Processor Freeze and System Freeze should only be done by, or under the direction of, your service provider. When System Freeze is enabled and one freeze-enabled processor halts, all other freeze-enabled processors in the system also halt. When enabled, ServerNet disruptions such as cable replacement or CRU/FRU insertion can generate a system freeze. If System Freeze is enabled, disable it before performing a service operation.

The system and processors are freeze-disabled by default. You can check the current state through the following attributes:

- **System Freeze**—In the OSM Low-Level Link, under the System object, the System Freeze attribute indicates whether System Freeze is currently Enabled or Disabled.
- **Processor Freeze**—Check either:
 - In the OSM Low-Level Link, under each Processor object, the Processor Freeze attribute indicates whether Processor Freeze is currently Enabled or Disabled for that processor.
 - In the OSM Service Connection, under each Logical Processor object, the Processor Freeze State attribute indicates whether Processor Freeze is currently Enabled or Disabled for that processor.

To enable or disable System Freeze, use Enable System Freeze action, located under the System object in the OSM Low-Level Link. After a System Freeze action is performed, the System Freeze attribute is automatically set back to Disabled.

To enable or disable Processor Freeze, you can use either the OSM Low-Level Link or the OSM Service Connection:

- In the OSM Low-Level Link, use the Enable Freeze or Disable Freeze actions, located under the Processor object.
- In the OSM Service Connection, use the Enable Processor Freeze or Disable Processor Freeze actions, located under the Logical Processor object.

Freezing the System and Freeze-Enabled Processors

In the OSM Low-Level Link, under the System object, perform the System Freeze action. This action halts all freeze-enabled processors in the system.

Confirm that the action success Processor Freeze State for each processor is now Enabled. In addition to the attribute values described earlier, in the LLL Processor Status dialog box, each processor should now display an “F” appears next to its name

Dumping a Processor to Disk

Dump options for NonStop NS-series servers are different than for NonStop S-series servers. While dumping to tape is not option for NS-series, there are many new options for dumping an entire processor or just the processor element (PE) needed.

For automatic dumping and reloading of halted processors, use the HP Tandem Failure Data System (TFDS). To dump automatically, TFDS must be configured on the system before the halt occurs. However, you can also bring up TFDS following a halt and use the TFDS`COM` command `ANALYZE CPU` to obtain information about the failure.

If you did not have TFDS configured to take the processor dump, you can use the `RCVDUMP` utility to take the dump.

- If your service provider determines that a processor halt is divergence-related, you might be directed to dump the entire processor before reloading it. In this case, use the `RCVDUMP` command, as follows:
 - Use `BLADE ALL` parameter option.
 - Do not specify `ONLINE` or `PARALLEL` parameters.

See [Using RCVDUMP to Dump a Processor to Disk](#) on page 9-17.

- If your service provider determines that a processor halt is not divergence-related, you might be directed to reload the processor while excluding the PE for one Blade Element, which is then dumped before being reintegrated. In this case, perform the reload (see [Reloading a Single Processor on a Running Server](#) on page 9-10), then use the `RCVDUMP` command as follows:
 - If more than one Blade Element is in the Stopped state, use the `BLADE` parameter and specify the *bladeId* (*A*, *B*, or *C*) of the PE to be dumped. (If only one Blade Element is in the Stopped state, it is not necessary to use the `BLADE` parameter.)
 - Specify the `PARALLEL` parameter.

See [Using RCVDUMP to Dump a Processor to Disk](#) on page 9-17.

- If a dump is to be taken following a system load, as described in [Performing a System Load From a Specific Processor](#) on page 15-11, options for taking dumps include:
 - After the processor (0 or 1) has been loaded with a Blade Element omitted, use `RCVDUMP` with the `PARALLEL` option.
 - You can dump any of the remaining processors either by dumping the entire processor before reloading them (use `RCVDUMP` without specifying the `ONLINE` or `PARALLEL` options), or reload with a Blade Element omitted, then dump that Blade Element using `RCVDUMP` with the `PARALLEL` option.
- If you are directed to dump a single PE that is running, use the `RCVDUMP` command, as follows:
 - Specify the `ONLINE` parameter.
 - Do not specify the Blade ID (`RCVDUMP` will choose the first running PE).

See [Using RCVDUMP to Dump a Processor to Disk](#) on page 9-17.

Before You Begin

- You must have a second processor connected to a terminal or workstation with a running command interpreter. The processor in which the TACL command interpreter is running performs the dump.
- If *dumpfile* already exists, it must be empty. (Its end-of-file pointer, or EOF, must be zero.)
- You must not prime or reset the processor before performing a processor dump.

To prepare for a disk dump:

1. Verify that a disk is available with enough space to store the dump.

A processor dump requires 256 extents. Each extent should equal slightly more than 1/256 the size of the processor memory. For example, for a processor with 256 megabytes of main memory, you need 256 extents of at least 512 pages each.

2. To empty an existing *dumpfile*:

```
> FUP PURGEDATA dumpfile
```

A processor dump can be made to one, two, or three files, with file names ending with an A, B, or C to designate which Blade Element the processor element was dumped from.

Using RCVDUMP to Dump a Processor to Disk

The RCVDUMP command has several new parameters for Integrity NonStop systems. See [Dumping a Processor to Disk](#) on page 9-15 for how they might be used for some possible dump scenarios.

The *parallel* parameter enables you to dump and reload individual physical processors without affecting the continuing operation of the other processor elements in that logical processor grouping.

To dump a processor to disk on a running system:

1. Log on to a TACL session as the super ID (255,255).
2. At a TACL prompt, run the RCVDUMP utility, choosing the parameters and options appropriate for your scenario:

```
RCVDUMP filename, cpuNum [, BLADE bladeId]  
[, START startAddress][, END endAddress]  
[, ONLINE | PARALLEL]
```

filename

is the name of the disk file to which the dump is to be written.

cpuNum

is the number of the logical processor from which a processor element is to be dumped. Specify *cpuNum* as an integer in the range from 0 through 15.

BLADE bladeId

is the identification of the Blade Element from which the processor element is to be dumped. Valid values are *A* or *B* or *C* or *ALL*. Note that *ALL* may not be used with the parallel method of dumping.

START n...

is the byte address where the dump will start. The default value is 0.

END n...

is the byte address where the dump will stop. Using a value of -1 is the same as specifying the end of memory. The default value is -1.

ONLINE

If this option is specified, a dump can be taken of a processor while it is running. You may use either *PARALLEL* or *ONLINE* but not both.

PARALLEL

If this option is specified, a dump may be taken of a single processor element while the other PEs in that logical processor are reloaded and continue normal operations. You may use either *PARALLEL* or *ONLINE* but not both.

For more information, see the *HP NonStop TACL Reference Manual*.

3. Monitor the dump to make sure that it finishes successfully:

a. Wait for this message to appear:

```
CPU n has been dumped to dumpfile
```

b. Check the size of *dumpfile* to verify that the end-of-file pointer (EOF) is not equal to zero.

```
> FUP INFO dumpfile
```

When a processor is dumped to disk, the RCVDUMP utility begins copying the dump in a compressed format from the specified processor into a disk file called *dumpfile*. If *dumpfile* does not exist, the RCVDUMP utility creates it. As the dump proceeds, the status of the processor being dumped changes in the Processor Status dialog box to indicate that a dump is in progress. When the dump is finished, the status of the selected processor in the Processor Status dialog box changes to indicate the completion of the dump.

Blade Element Reintegration

Whether TFDS or RCVDUMP was used to dump the PE for an individual Blade Element, reintegration should take place automatically upon completion of the dump.

The OSM Service Connection has a Reintegrate PE action, located under the Processor Components object, in case the automatic reintegration fails.

Troubleshooting and Recovery Operations for Disk Dumps

If a message indicates that the dump was not successful, repeat [Dumping a Processor to Disk](#) on page 9-13 using the other ServerNet fabric.

If a halt code for the selected processor appears in the Processor Status dialog box of the OSM Low-Level Link, look it up in the *Processor Halt Codes Manual* for further information about the cause of failure and the appropriate recovery procedure.

Backing Up a Processor Dump to Tape

Back up a processor dump to tape from the compressed disk file generated by the TACL RECEIVEDUMP command (or the RCVDUMP utility):

1. Follow the instructions in [Dumping a Processor to Disk](#) on page 9-15.
2. Use the BACKUP utility to copy the processor dump to tape, as described in [Section 10, Disk Drives: Monitoring and Recovery](#).
3. See [Submitting Information to Your Service Provider](#) on page 9-19.

Replacing Processor Memory

Processor memory is field-replaceable for all Integrity NonStop systems. Call your service provider. If memory units cannot be replaced, the memory board must be replaced.

If the Service State attribute of a Blade Element object is not OK (see [Monitoring Processor Performance Using ViewSys](#) on page 9-7), the memory board might need to be replaced. Contact your service provider.

Replacing the Processor Board and Processor Entity

Processor boards and entities are field-replaceable for all Integrity NonStop systems. Call your service provider.

Submitting Information to Your Service Provider

To help with the analysis of a processor dump, submit a backup tape of other system configuration and operations files and some additional information.

- [Submitting Tapes of Processor Dumps](#) on page 9-20
- [Submitting Tapes of Configuration and Operations Files](#) on page 9-20
- [Additional Information Required by Your Service Provider](#) on page 9-21

Submitting Tapes of Processor Dumps

Use a separate tape for each processor dump. For each tape you submit, record:

- The notation BACKUP to indicate a disk dump that has been backed up to tape.
- The file names of the files on tape (if the BACKUP, LISTALL command has been used)

Submitting Tapes of Configuration and Operations Files

To ensure that a processor dump is usable by your service provider, place the files listed in [Table 9-1](#) on one backup tape using the BACKUP utility. Contact your service provider for information about any other files they might need.

Table 9-1. Other Files to Submit to Your Service Provider

File	Description
\$SYSTEM.ZSYSCONF.CONFIG	System configuration database
\$SYSTEM.SYS _{nn} .CONFTEXT	System configuration file
\$SYSTEM.SYS _{nn} .CONFLIST	System generation program output file
\$SYSTEM.ZLOG _{nn}	EMS event log (\$0 operator log files)
All files located in the \$SYSTEM.ZSERVICE subvolume	Service event log (\$ZLOG files)

To back up configuration and operations files:

1. For this backup operation, use any tape drive that is in a STARTED state and a READY substate. To determine the names and current states of the tape drives on a system:

```
> SCF STATUS TAPE $*
```

2. To back up all the configuration and operations files to tape, use the BACKUP utility. For example:

```
>BACKUP $tape, (CPU0,$SYSTEM.SYS00.CONFTEXT,&
$SYSTEM.SYS00.CONFLIST,$SYSTEM.ZSYSCONF.CONFIG,&
$SYSTEM.ZLOG00.*, $SYSTEM.ZSERVICE.*),LISTALL,OPEN,VERIFYREEL
```

For more information on the BACKUP command and options, see the *Guardian Disk and Tape Utilities Reference Manual*.

Additional Information Required by Your Service Provider

In addition to the tapes previously discussed, submit the information listed in [Table 9-2](#) to your service provider.

Table 9-2. Additional Processor Dump Information for Your Service Provider

Name of HP branch office	_____
Your company name	_____
System number	_____
The processor numbers of the processors that were dumped, along with the letter designation of the PEs dumped.	_____
The date that the processor dump was done	_____
The RVU you are using	_____

You should also provide:

- A list of any software product revisions (SPRs) you have installed since installing the RVU.
 - A list of any customer-written privileged programs running on your system and explanations of what they do.
 - The reason for the processor dump. If you performed a processor dump because a processor halted, include the halt code and the frequency of the halts. The halt code and other information is displayed in the Processor Status dialog box of the OSM Low-Level Link.
 - Any particular circumstances that you think contributed to the problem. A brief description of the problem in an EDIT file and a short program that duplicates or illustrates the problem (if possible) would be helpful.
 - Any OSM status messages reported.
-

Related Reading

For more information about tools used to monitor and perform recovery operations on processors, refer to the documentation listed in [Table 9-3](#).

Table 9-3. Related Reading for Monitoring and Recovery Operations on Processors

For Information About	Tool	See
Recovery operations for processes	TACL	<i>Guardian User's Guide</i> <i>TACL Reference Manual</i>
Monitoring or recovery operations on processors	OSM	OSM online help <i>OSM Service Connection User's Guide</i>
Replacing a memory unit	Usually requires service provider	Replace Blade Element Procedure (Available in the OSM Service Connection, accessed by performing the Blade Element Replace action.)
Monitoring processor performance	ViewSys	<i>ViewSys User's Guide</i>
Recovery operations for processor halt	TFDS	<i>Tandem Failure Data System (TFDS) Manual</i>

10

Disk Drives: Monitoring and Recovery

[When to Use This Section](#) on page 10-1

[Overview of Disk Drives](#) on page 10-2

[Internal SCSI Disk Drives](#) on page 10-2

[M8xxx Fibre Channel Disk Drives](#) on page 10-3

[Enterprise Storage System \(ESS\) Disks](#) on page 10-3

[Monitoring Disk Drives](#) on page 10-4

[Monitoring Disk Drives With OSM](#) on page 10-4

[Monitoring Disk Drives With SCF](#) on page 10-5

[Monitoring the State of Disk Drives](#) on page 10-9

[Monitoring the Use of Space on a Disk Volume](#) on page 10-9

[Monitoring the Size of Database Files](#) on page 10-9

[Monitoring Disk Configuration and Performance](#) on page 10-10

[Identifying Disk Drive Problems](#) on page 10-11

[Identifying Disk Drive Problems](#) on page 10-11

[Recovery Operations for Disk Drives](#) on page 10-12

[Common Recovery Operations for Disk Drives](#) on page 10-12

[Recovery Operations for a Down Disk or Down Disk Path](#) on page 10-14

[Recovery Operations for a Nearly Full Database File](#) on page 10-15

[Related Reading](#) on page 10-15

When to Use This Section

Use this section to monitor M8xxx Fibre Channel and internal SCSI disk drives, and to recover from common disk problems.

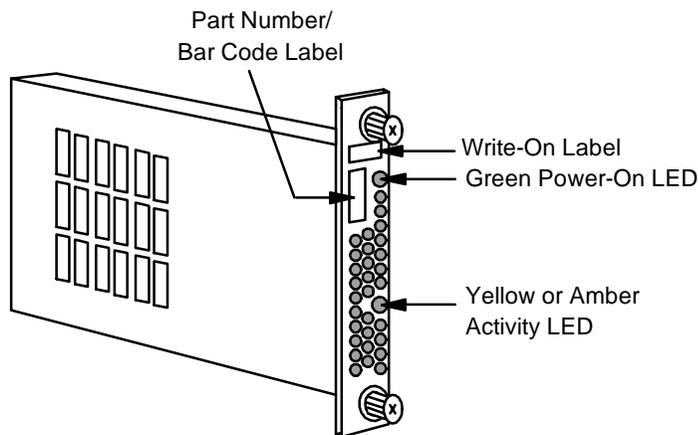
Overview of Disk Drives

The Integrity NonStop NS-series server supports three types of disk drives:

- [Internal SCSI Disk Drives](#)
- [M8xxx Fibre Channel Disk Drives](#)
- [Enterprise Storage System \(ESS\) Disks](#)

Internal SCSI Disk Drives

Internal SCSI disk drives are installed in NonStop S-series I/O enclosures.



CDT 106CDD

These disk drives are Class-1 CRUs.

Any physical action on a CRU, including installing and replacing disks, can be performed by customers. However, depending on the class of CRU, training in replacement techniques might be recommended.

For information about

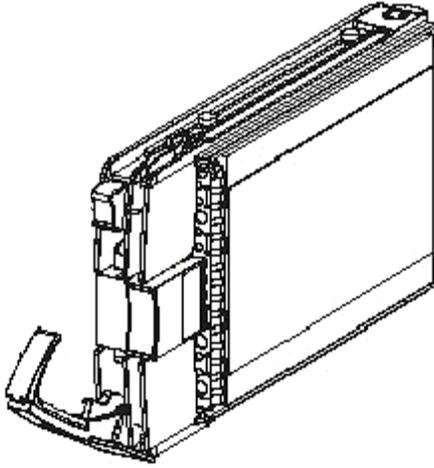
Internal SCSI disk specifications
Internal SCSI disk commands
Classes of CRUs

See

NonStop S-Series Planning and Configuration Guide
SCF Reference Manual for the Storage Subsystem
NonStop S-Series Planning and Configuration Guide

M8xxx Fibre Channel Disk Drives

M8xxx Fibre Channel disk drives are installed in disk drive enclosures. A single disk drive is shown:



VST601.vsd

Fibre Channel disk drives are field-replaceable units (FRUs). Any physical action on a FRU, including installation and replacement, must be performed only by a qualified HP service provider.

For information about

M8xxx Fibre Channel disk specifications

M8xxx disk commands

FRUs and CRUs

See

Integrity NonStop NS-Series Planning Guide

SCF Reference Manual for the Storage Subsystem

NonStop S-Series Planning and Configuration Guide

Enterprise Storage System (ESS) Disks

The Enterprise Storage System (ESS) is any of several models of HP StorageWorks Disk Arrays. These arrays are a collection of magnetic disks, their controllers, and the disk cache in one or more standalone cabinets. The disks are configured with an attached console.

ESS disk drives are maintained and serviced only by qualified service providers who have completed courses in ESS management.

For information about

ESS disk drives

ESS models supported on Integrity NonStop NS-series systems

See

HP XP StorageWorks documentation

Your HP representative

Monitoring Disk Drives

For information about

Tools for monitoring status, space use, configuration, and performance of disk drives

Displaying EMS events generated by storage devices and subsystems

Monitoring with OSM

Monitoring with SCF

See

[Appendix B, Tools and Utilities for Operations](#)

[Section 4, Monitoring EMS Event Messages](#)

[Monitoring Disk Drives With OSM](#) on page 10-4

[Monitoring Disk Drives With SCF](#) on page 10-5

Monitoring Disk Drives With OSM

Task

Monitor the status of disk drives

Inventory the entire system, including disk drives

Use:

- OSM Service Connection
- OSM Event Viewer
- OSM Inventory View

Determine:

- Service state
- Primary path state
- Secondary path state

Learn possible values of primary and backup path state attributes for disk drives and disk paths

See

- OSM Service Connection
- OSM Event Viewer

OSM Inventory View

You can save this view as a file in Excel.

OSM Online Help

[Using OSM to Monitor the System](#) on page 3-7

[Monitoring the State of Disk Drives](#) on page 10-9

Monitoring Disk Drives With SCF

This subsection explains how to list disk volumes and determine their status.

1. List the status of all magnetic disk volumes on your system, issue this command from SCF:

```
> STATUS DISK $*, SUB MAGNETIC
```

1-> STATUS DISK \$*, SUB MAGNETIC						
STORAGE - Status DISK \COMM.\$SYSTEM						
LDev	Primary	Backup	Mirror	MirrorBackup	Primary PID	Backup PID
6	*STARTED	STARTED	*STARTED	STARTED	0,257	1,257
STORAGE - Status DISK \COMM.\$VIRCFG						
LDev	Primary	Backup	Mirror	MirrorBackup	Primary PID	Backup PID
146	*STARTED	STARTED	*STARTED	STARTED	2,288	3,267
STORAGE - Status DISK \COMM.\$WORK2						
LDev	Primary	Backup	Mirror	MirrorBackup	Primary PID	Backup PID
140	*STARTED	STARTED	*STARTED	STARTED	5,278	4,273
STORAGE - Status DISK \COMM.\$WEB02						
LDev	Primary	Backup	Mirror	MirrorBackup	Primary PID	Backup PID
143	*STARTED	STARTED	*STARTED	STARTED	2,289	3,266
STORAGE - Status DISK \COMM.\$ROOT						
LDev	Primary	Backup	Mirror	MirrorBackup	Primary PID	Backup PID
190	*STARTED	STARTED	*STARTED	STARTED	3,268	2,287
STORAGE - Status DISK \COMM.\$P1D02						
LDev	Primary	Backup	Mirror	MirrorBackup	Primary PID	Backup PID
247	*STARTED	STARTED	*STARTED	STARTED	4,268	5,276
STORAGE - Status DISK \COMM.\$P1D03						
LDev	Primary	Backup	Mirror	MirrorBackup	Primary PID	Backup PID
246	*STARTED	STARTED	*STARTED	STARTED	4,269	5,282

2. Get information about a disk with SCF STATUS DISK, DETAIL. For example:

-> STATUS DISK \$DATA09, DETAIL

The output from this example shows that \$DATA09 is in the STOPPED state, HARDDOWN substate.

```
STORAGE - Detailed Status DISK \SHARK.$DATA09
Disk Path Information:
  LDev  Path                Status      State      Substate   Primary  Backup
        PID                PID
  92    PRIMARY              INACTIVE   STOPPED   HARDDOWN   2,266   3,266
  92    BACKUP                 INACTIVE   STOPPED   HARDDOWN   2,266   3,266
  92    MIRROR                 INACTIVE   STOPPED   HARDDOWN   2,266   3,266
  92    MIRROR-BACKUP         INACTIVE   STOPPED   HARDDOWN   2,266   3,266

General Disk Information:
Device Type..... 3                Device Subtype..... 53
Primary Drive Type....           Mirror Drive Type.....
Physical Record Size.. 4096        Priority..... 220
Library File.....
Program File..... $SYSTEM.SYS00.TSYS2P2
Protection..... MIRRORED

Hardware Information:
  Path                Location      Power      Physical Status
        (group,module,slot)
  PRIMARY             EXTERNAL     DUAL       PRESENT
  MIRROR              EXTERNAL     NONE       ABSENT
Total Errors = 0      Total Warnings = 0
```

3. See [Recovery Operations for a Down Disk or Down Disk Path](#) on page 10-14.

Examples

To display the status of the disk \$DATA01:

-> STATUS \$DATA01

```
34-> STATUS $DATA01
STORAGE - Status DISK \SHARK.$DATA01
LDev  Primary  Backup  Mirror  MirrorBackup  Primary  Backup
        PID                PID
  63  *STARTED  STARTED  *STARTED  STARTED        0,267   1,266
```

To display the status of the mirror disk of the volume \$DATA02:

-> STATUS \$DATA02-M

```
47-> STATUS DISK $DATA02-M
STORAGE - Status DISK \SHARK.$DATA02-M
LDev  Path                PathStatus  State      SubState   Primary  Backup
        PID                PID
  62  MIRROR              INACTIVE   STOPPED   HARDDOWN   0,268   1,265
```

To display the status of all disks:

-> STATUS DISK \$*

```

1-> STATUS DISK $*
STORAGE - Status DISK \COMM.$SYSTEM
LDev  Primary  Backup  Mirror  MirrorBackup  Primary  Backup
      PID    PID    PID    PID    PID    PID    PID
   6  *STARTED  STARTED  *STARTED  STARTED      0,257   1,257

STORAGE - Status VIRTUAL DISK \COMM.$VIEWPT
LDev  State    Primary  Backup  Type  Subtype
      PID    PID    PID
  147  STARTED   9,22   8,53   3    36

STORAGE - Status VIRTUAL DISK \COMM.$WANA
LDev  State    Primary  Backup  Type  Subtype
      PID    PID    PID
  145  STARTED   8,77   9,56   3    36

STORAGE - Status VIRTUAL DISK \COMM.$WEB
LDev  State    Primary  Backup  Type  Subtype
      PID    PID    PID
  144  STARTED   9,29   8,48   3    36

STORAGE - Status VIRTUAL DISK \COMM.$WEBVPT
LDev  State    Primary  Backup  Type  Subtype
      PID    PID    PID
  142  STARTED   9,26   8,47   3    36

STORAGE - Status VIRTUAL DISK \COMM.$WIPRO
LDev  State    Primary  Backup  Type  Subtype
      PID    PID    PID
  141  STARTED   9,27   8,51   3    36

STORAGE - Status VIRTUAL DISK \COMM.$ZERO
LDev  State    Primary  Backup  Type  Subtype
      PID    PID    PID
  133  STARTED   8,78   9,57   3    36

STORAGE - Status VIRTUAL DISK \COMM.$ZIMBU
LDev  State    Primary  Backup  Type  Subtype
      PID    PID    PID
  115  STARTED   9,28   8,52   3    36

STORAGE - Status DISK \COMM.$VIRCFG
LDev  Primary  Backup  Mirror  MirrorBackup  Primary  Backup
      PID    PID    PID    PID    PID    PID
  146  *STARTED  STARTED  *STARTED  STARTED      2,288   3,267
    
```

To display the detailed status of the disk \$DATA01:

-> STATUS \$DATA01, DETAIL

```
STORAGE - Detailed Status DISK \SHARK.$DATA01
```

Disk Path Information:

LDev	Path	PathStatus	State	SubState	Primary PID	Backup PID
63	PRIMARY	ACTIVE	STARTED		0,267	1,266
63	BACKUP	INACTIVE	STARTED		0,267	1,266
63	MIRROR	ACTIVE	STARTED		0,267	1,266
63	MIRROR-BACKUP	INACTIVE	STARTED		0,267	1,266

General Disk Information:

```
Device Type..... 3           Device Subtype..... 53
Primary Drive Type.... BF0365   Mirror Drive Type..... BF0365
Physical Record Size.. 4096     Priority..... 220
Library File.....
Program File..... \SHARK.$SYSTEM.SYS00.TSYS02
Protection..... MIRRORED
```

Usage Information:

```
Capacity (MB)..... 36419.03   Free Space (MB).....33671.23 (92.45%
Free Extents..... 14         Largest Free Extent (MB). 33516.31
```

Hardware Information:

Path	Location (group,module,slot)	Power	Physical Status
PRIMARY	EXTERNAL	DUAL	PRESENT
MIRROR	EXTERNAL	DUAL	PRESENT

To display status of all paths for \$DATA00:

-> STATUS DISK \$DATA00-*

```
STORAGE - Status DISK \ALM171.$DATA00-*
```

LDev	Path	PathStatus	State	SubState	Primary PID	Backup PID
6	PRIMARY	ACTIVE	STARTED		0,10	1,10
6	BACKUP	INACTIVE	STARTED		0,10	1,10
6	MIRROR	ACTIVE	STARTING	REVIVE	0,10	1,10
6	MIRROR-BACKUP	INACTIVE	STARTING	REVIVE	0,10	1,10

The output from this example indicates that \$DATA00:

- Is a mirrored volume (primary and mirror paths)
- Has a mirror disk that is being revived (SubState REVIVE)

- LDev Logical device number
- Path Disk path assignment
- PathStatus Status of the disk path; whether that path is the current path (ACTIVE) or not (INACTIVE)
- State Current SCF state of the disk path
- SubState Current SCF substate of the disk path
- Primary PID Primary processor number and process identification number (PIN) of the specified device
- Backup PID Backup processor number and PIN of the specified device

Monitoring the State of Disk Drives

Each disk drive can have two paths, the primary and the backup. (Each M8xxx disk drive is forced to have two paths.) The two path states are represented separately.

Table 10-1. Primary and Backup Path States for Disk Drives

Path State	Description
Degraded	This path of this disk drive has a state other than Up.
Down	The disk volume or disk path is not logically accessible.
Exclusive	Exclusive ownership has been declared for this disk volume. The disk is not accessible to other users.
Executing Diagnostics	The processor is performing diagnostics.
Format in Progress	A disk-format operation is in progress.
Hard Down	The volume or path was put in this state by the SCF ABORT DISK command, or cannot be accessed because of a hardware error.
Inaccessible	The disk cannot be accessed.
Not Configured	The component is not configured.
Revive	A mirrored disk is being updated.
Special	Only maintenance-type I/O tasks can be performed on the disk.
Unknown	The path state is unknown. The disk might not be responding.
Up	The disk volume or disk path is logically accessible.

Monitoring the Use of Space on a Disk Volume

The Disk Space Analysis Program (DSAP) provides information on disk capacity, free-space fragments, and page allocation. To check for bad sectors, you can use SCF:

```
-> INFO DISK $*, BAD, SEL started, sub magnetic
```

```
Bad Sectors Information $DATA14 Primary:
No bad sectors found.
```

You can also use DSAP at a TACL prompt:

```
> DSAP $*
```

Monitoring the Size of Database Files

To check file size:

```
> FUP INFO filename, DETAIL
```

A report similar to this one is sent to your home terminal:

```

$DATA.FILES.FILEA  10 Jul 1993, 14:05
  ENSCRIBE
  TYPE U
  CODE 100
  EXT ( 224 PAGES, 14 PAGES )
  ODDUNSTR
  MAXEXTENTS 370
  BUFFERSIZE 4096
  OWNER 8,255
  SECURITY (RWEP): NUNU, LICENSED
  DATA MODIF:  10 Jul 1994, 14:04
  CREATION DATE:  10 Jan 1994, 14:04
  LAST OPEN:  10 Jul 1994, 14:04
  EOF 267022 (58.2% USED)
  FILE LABEL: 822 (20.2% USED)
  EXTENTS ALLOCATED: 10

```

This report shows that FILEA is 58.2% full. If a database file is 90% full or more, see [Recovery Operations for a Nearly Full Database File](#) on page 10-15.

Example

To check the size of the file DATA1.MEMOS:

```
> FUP INFO DATA1.MEMOS, DETAIL
```

```

$DATA.DATA1.MEMOS  12 Jul 1994, 14:05
  ENSCRIBE
  TYPE U
  CODE 101
  EXT ( 2 PAGES, 2 PAGES )
  ODDUNSTR
  MAXEXTENTS 16
  BUFFERSIZE 4096
  OWNER 8,255
  SECURITY (RWEP): NUNU
  DATA MODIF:  12 Jul 1994, 14:04
  CREATION DATE:  12 Jan 1994, 14:04
  LAST OPEN:  12 Jul 1994, 14:04
  EOF 567022 (88.2% USED)
  FILE LABEL: 775 (31.6% USED)
  EXTENTS ALLOCATED: 10

```

Monitoring Disk Configuration and Performance

For information about	See
Checking configuration information for disk devices	<i>SCF Reference Manual for the Storage Subsystem</i>
Monitoring disk block and cache statistics	<i>SCF Reference Manual for the Storage Subsystem</i>
Examining system performance data with Measure	<i>Measure User's Guide</i>

Identifying Disk Drive Problems

For recovery operations, refer to [Recovery Operations for Disk Drives](#) on page 10-12.

Table 10-2. Possible Causes of Common Disk Drive Problems

Problems	Possible Symptoms
Disk is full or does not have enough space.	Error 43 (unable to obtain disk space for file extent) occurs. If the disk is full, an application might go down.
Disk free space is fragmented.	
One disk in a mirrored pair is down.	The storage subsystem generates an event message but the application continues to run.
An unmirrored disk is down, or both disks in a mirrored pair are down.	Users report access problems, applications go down, and the storage subsystem generates event messages.
Performance problems occur due to path switches or a cache size that is too small.	Users report poor application performance. Output from SCF INFO DISK, BAD command indicates unspared defective sectors. Intm-errors-exceeded message Slow-IOs-threshold-exceeded message
Defective tracks or sectors exist.	
Disk errors exceed a certain limit.	
Slow I/O operations exceed a certain limit.	

Internal SCSI Disk Drives

The most common disk drive problems on a NonStop S-series server include:

- Space problems such as full disks or free-space fragmentation
- Stopped disks
- Performance problems
- Defective tracks or sectors

M8xxx Fibre Channel Disk Drives

The most common disk problems on an Integrity NonStop NS-series server are intm-errors-exceeded and slow-IOs-threshold-exceeded errors on the Fibre Channel loop.

Such errors are often normal. However, if they cause problems on a Fibre Channel loop, power the affected disk down and up again. This procedure can solve the problem temporarily.

Unless you are a qualified service provider, you cannot perform any physical actions on disk drives. However, operators can use OSM and SCF commands.

Recovery Operations for Disk Drives

These SCF commands control DISK objects:

Command	Description
ABORT	Terminates the operation of a disk drive immediately, leaving it in the STOPPED state, HARDDOWN substate.
ALTER	Changes attribute values for a storage device.
BYPASS	Bypasses one or more disks in a Fibre Channel disk-drive enclosure.
CONTROL	Issues disk-specific commands.
PRIMARY	For a disk drive, causes the backup processor to become the primary processor and the primary processor to become the backup processor.
RENAME	Changes the name of a disk drive.
RESET	Puts a disk drive in a state from which it can be restarted.
START	Initiates the operation of a disk drive.
STOP	Terminates the operation of a disk drive in a normal manner.
SWITCH	Switches paths to a disk drive.

For more information, see the *SCF Reference Manual for the Storage Subsystem*.

Table 10-3. Common Recovery Operations for Disk Drives (page 1 of 2)

Problem	Recovery
Free-space fragmentation	Use the Disk Compression Program (DCOM) to consolidate disk space. See Disk Compression Program (DCOM) on page B-2.
Disk full	<ol style="list-style-type: none"> 1. Use DSAP to identify large, old, and little used files. 2. If you are authorized: <ul style="list-style-type: none"> ● Use the BACKUP utility to back up these disk files to tape and then purge them from the disk. Do not purge important system files. ● Move files to another disk. Do not move important system files. ● Ask users to purge files. <p>For more information about these utilities, see BACKUP on page B-2 and Disk Space Analysis Program (DSAP) on page B-2.</p>
Down disk or disk path	Recovery Operations for a Down Disk or Down Disk Path on page 10-14
Defective sectors	<p>If you are authorized, use the SCF CONTROL DISK, SPARE command to spare defective sectors. For information on reinitializing the disk drive, see the <i>SCF Reference Manual for the Storage Subsystem</i>.</p> <p>Disks come formatted from HP. No disk format utility is available. Return any disk that requires formatting to HP.</p>

Table 10-3. Common Recovery Operations for Disk Drives (page 2 of 2)

Problem	Recovery
Unspared defective sectors	<p>To check for unspared defective sectors with SCF:</p> <pre>-> INFO DISK \$*, BAD, SEL started, sub magnetic</pre> <p>To check for unspared defective sectors with DSAP at a TACL prompt:</p> <pre>> DSAP \$*</pre> <ul style="list-style-type: none"> ● Recovery for DSAP is not needed. ● Recovery for DCOM: use the SCF INFO DISK, BAD command on the affected disk to obtain the bad sector address. Before restarting DCOM, perform the CONTROL DISK, SPARE command. <p>For more information, see the <i>Guardian Disk and Tape Utilities Manual</i>.</p>
Nearly full database file	<p>Recovery Operations for a Nearly Full Database File on page 10-15</p>
Performance problems	<p>Performance problems can have various causes, including path switches or a cache size that is too small.</p> <p>For information about disk load balancing and increasing cache size, see the <i>SCF Reference Manual for the Storage Subsystem</i>.</p>
Corrupt \$SYSTEM disk	<p>If both halves of your mirrored system volume become corrupted, use an alternate system disk if one is available. For how to create an alternate system disk, see the <i>Integrity NonStop NS-Series Planning Guide</i>.</p> <ul style="list-style-type: none"> ● For internal SCSI disk drives: if there is no alternate system disk and you cannot load from the CONFBASE file, you might be able to perform a tape load from a system image tape (SIT) to restore the system image files to the \$SYSTEM disk (SYS_{nn} and CSS_{nn} subvolumes). Then load that image into processor 0 or 1. <p>A tape load reinitializes the disk directory. The disk directory is overlaid with the directory from the tape. All files on that disk are destroyed. Perform a tape load only with the advice of the Global Customer Support Center or your service provider.</p> <ul style="list-style-type: none"> ● For M8xxx disk drives: you cannot perform a tape load from a SIT.
Failed disk drives	<ul style="list-style-type: none"> ● Internal SCSI disks: the Support and Service Library on page 1-12 describes replacing disk drives. ● M8xxx fibre channel disks: these disks are FRUs and can be serviced or replaced only by HP-trained service personnel.

Recovery Operations for a Down Disk or Down Disk Path

To restart a disk or disk path:

1. If a path is down due to a ServerNet fabric failure, determine the affected paths. From an SCF prompt:

```
-> STATUS DISK $*-* , SUB MAGNETIC
```

The output indicates:

- \$DATA06-M and \$DATA06-MB are stopped in the DOWN substate.
- \$WD8-M and \$WD8-MB are stopped in the HARDDOWN substate.
- \$DATA00-P and \$DATA00-B are stopped in the HARDDOWN substate.

STORAGE - Status DISK \ALPHA12.\$DATA06-*							
LDev	Path	Status	State	Substate	Primary PID	Backup PID	
116	PRIMARY	ACTIVE	STARTED		0,285	1,268	
116	BACKUP	INACTIVE	STARTED		0,285	1,268	
116	MIRROR	INACTIVE	STOPPED	DOWN	0,285	1,268	
116	MIRROR-BACKUP	INACTIVE	STOPPED	DOWN	0,285	1,268	
STORAGE - Status DISK \ALPHA12.\$WD8-*							
LDev	Path	Status	State	Substate	Primary PID	Backup PID	
96	PRIMARY	ACTIVE	STARTED		0,23	1,12	
96	BACKUP	INACTIVE	STARTED		0,23	1,12	
96	MIRROR	INACTIVE	STOPPED	HARDDOWN	0,23	1,12	
96	MIRROR-BACKUP	INACTIVE	STOPPED	HARDDOWN	0,23	1,12	
STORAGE - Status DISK \ALPHA12.\$DATA00-*							
LDev	Path	Status	State	Substate	Primary PID	Backup PID	
121	PRIMARY	INACTIVE	STOPPED	HARDDOWN	0,284	1,267	
121	BACKUP	INACTIVE	STOPPED	HARDDOWN	0,284	1,267	
121	MIRROR	ACTIVE	STARTED		0,284	1,267	
121	MIRROR-BACKUP	INACTIVE	STARTED		0,284	1,267	
Total Errors = 0		Total Warnings = 9					

2. Reset disk drives that are in the HARDDOWN substate. At an SCF prompt:

```
-> RESET DISK $volume
```

For example:

```
-> RESET DISK $WD8
```

Resetting a disk in the HARDDOWN substate places it in the DOWN substate.

3. Restart the disk. At an SCF prompt:

```
-> START DISK $volume
```

If the disk does not start, the disk might need replacement. If neither half of a mirrored volume starts, the database might need recovery. Contact your service provider.

Recovery Operations for a Nearly Full Database File

When a database file is 90 percent full or more, you can modify the file extents dynamically with FUP or perform other procedures according to your system policies.

Note. Allocating additional extents to any file causes that file to take up more disk space. Before you change the maximum allowable extents for any file, as shown in the next example, check your local procedures to determine whether this is the appropriate action for you to take.

To allocate additional extents to the file MEMOS:

```
> FUP
- ALTER MEMOS, MAXEXTENTS 20
- INFO MEMOS, DETAIL
```

A report such as this one is sent to your home terminal:

```
$DATA.DATA1.MEMOS 12 Jul 1993, 14:05
  ENSCRIBE
  TYPE U
  CODE 101
  EXT ( 2 PAGES, 2 PAGES )
  ODDUNSTR
  MAXEXTENTS 20
  BUFFERSIZE 4096
  OWNER 8,255
  SECURITY (RWEPR): NUNU
  DATA MODIF: 12 Jul 1993, 14:04
  CREATION DATE: 12 Jan 1993, 14:04
  LAST OPEN: 12 Jul 1993, 14:24
  EOF 567022 (78.5% USED)
  FILE LABEL: 649 (22.8% USED)
  EXTENTS ALLOCATED: 10
```

This report shows that the maximum extents allocated to this file have been increased to 20 and that the file MEMOS is now only 78.5% full.

Related Reading

For information about

Complete syntax, examples, and considerations for the SCF commands used in this section

Utilities such as:

- DCOM
- DSAP
- BACKUP
- RESTORE

Other operations procedures involving disk drives

See

SCF Reference Manual for the Storage Subsystem

Guardian Disk and Tape Utilities Reference Manual

Guardian User's Guide

11

Tape Drives: Monitoring and Recovery

[When to Use This Section](#) on page 11-1

[Overview of Tape Drives](#) on page 11-1

[Monitoring Tape Drives](#) on page 11-2

[Monitoring Tape Drive Status With OSM](#) on page 11-2

[Monitoring Tape Drive Status With SCF](#) on page 11-5

[Monitoring Tape Drive Status With MEDIACOM](#) on page 11-6

[Monitoring the Status of Labeled-Tape Operations](#) on page 11-7

[Identifying Tape Drive Problems](#) on page 11-7

[Recovery Operations for Tape Drives](#) on page 11-8

[Recovery Operations Using the OSM Service Connection](#) on page 11-8

[Recovery Operations Using SCF](#) on page 11-9

[Related Reading](#) on page 11-9

When to Use This Section

This section provides an overview of operating, monitoring, and recovery operations for tape drives attached to Integrity NonStop NS-series servers.

Overview of Tape Drives

A new generation of multimode fibre channel (MMF) peripherals is supported on Integrity NonStop systems. Tape drives with an MMF interface are connected directly to a fibre port on an FCSA in an IOAME or VIO enclosure. Some high voltage differential (HVD) SCSI drives are also supported in the NS optical environment. These drives are connected using an M8201 fibre channel to SCSI router. The M8201 converts the FCSA MMF connection into two SCSI buses. Currently, only optical port 0 is used on the M8201. The SCSI drives are connected using 68-pin to 68-pin HVD SCSI cables. These cables have a high density connector for the M8201.

Most drives supported on NonStop S-series can be connected to an Integrity NonStop NS16000 server through an IOMF2 in an IOMF enclosure. The drives can be interfaced in the same manner that is supported in S-series, either through an SNDA or a SCSI port on an IOMF2 board. OSM provides different views for drives that are connected through an FCSA versus an IOMF2. See [Monitoring Tape Drives](#) on page 11-2.

For further information, refer to the document on *Integrity NonStop NS-Series Supported Hardware* and the *NonStop NS16000 Hardware Installation Manual*.

Monitoring Tape Drives

This section describes the various methods of tape drives, which include:

- OSM Service Connection
- SCF
- Use MEDIACOM to monitor the use of tape drives and to write tape labels.

Monitoring Tape Drive Status With OSM

To check the status of all tape drives on your system:

1. Log on to the OSM Service Connection.
2. In the tree pane, expand the system object and check the Tape Collection object. A yellow arrow displayed over the Tape Collection object (see [Figure 11-1](#)) indicates that a problem exists with one or more of the tape drives connected to the system.
3. Expand the Tape Collection object and select the tape drive displaying a red or yellow triangular symbol over the tape drive object or bell-shaped symbol next to the object. For an example of an FCSA-connected tape drive, see [Figure 11-1](#); for an IOMF2-connected tape drive, see [Figure 11-2](#).
 - If a red or yellow triangular symbol is displayed over the tape drive object, check the Attributes tab for the specific attribute reporting a degraded value.
 - If a bell-shaped symbol is displayed next to the object, select the Alarms tab, click to select the alarm, then right-click and select Details to get more information about the alarm.

Figure 11-1. OSM: Monitoring Tape Drives Connected to an FCSA

The screenshot shows the OSM Service Connection interface. On the left is a tree view under 'System \OSMQA4' containing 'Group 110', 'Fabric Group 100', 'Monitored Service LAN Device', 'Processor Complex 400', and 'Tape Collection'. Under 'Tape Collection', several tape drives are listed: '\$LT02' (with a yellow warning icon), '\$SDLT', '\$VTAPE', '\$VTAPE01', '\$VTAPE03', and '\$VTAPE06'. On the right, the 'attributes' tab is selected for 'Tape Drive \$LT02'. The attributes are categorized into Logical, Associated Device, and Controller Path.

Tape Drive \$LT02	
Logical	
Device State	Hard Down
Firmware Version	
Logical Device Number	199
Product Id	Unknown
Ready Status	Not Ready
Tape Type	Unknown
Associated Device	
Fibre Channel Router	FC Router fc
Controller Path	
Configured FSAC Location	FCSA.SAC-1.GRP-110.MOD-3.SLOT-2
FSAC	FCSA.FSAC (110.3.2.1)
LUN	1
Port Name	0x100000E0022312E1

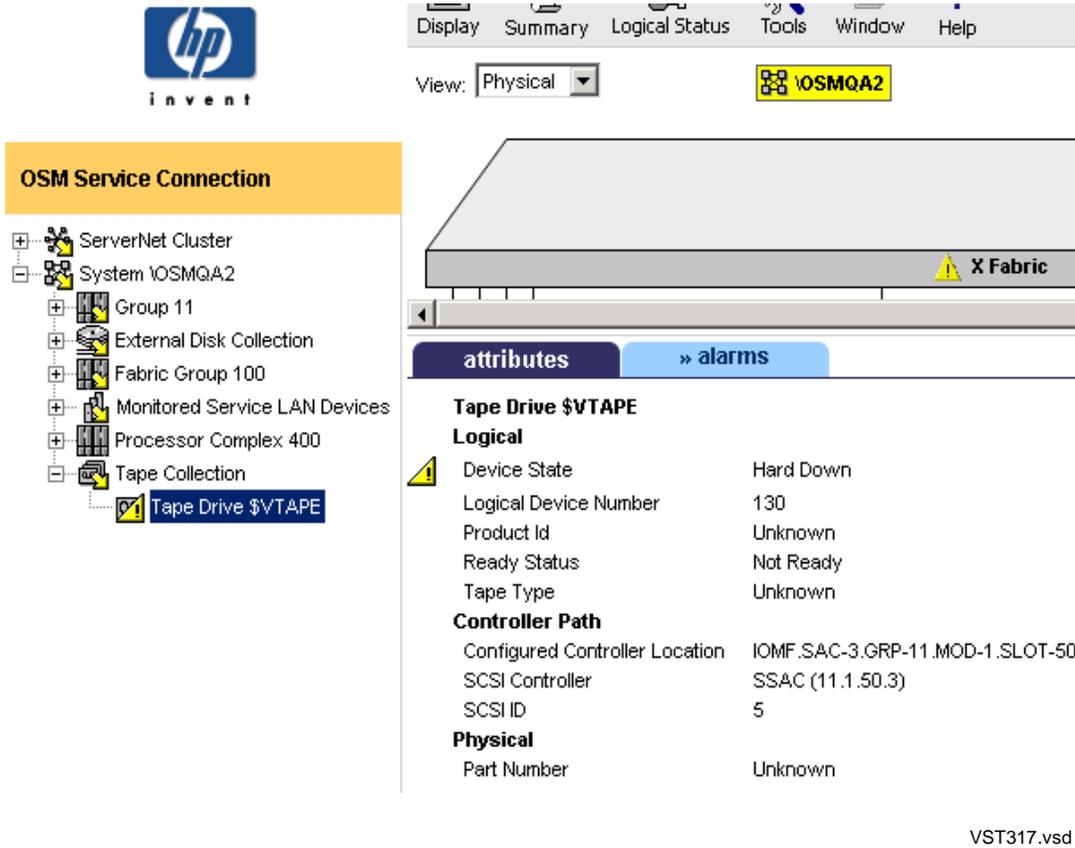
VST316.vsd

4. If an alarm or degraded conditions exist, the tape drive probably requires either:
 - Operator intervention. For more information, see [Recovery Operations for Tape Drives](#) on page 11-8.
 - Service or replacement. Contact your service provider and refer to the [Support and Service Library](#) on page 1-12 for the replacement procedure.

An alternative (to the method described above) for monitoring tape drives in OSM is to use the Multi-Resource Actions dialog box (available from the Display menu). In this dialog box, select the Tape Drive object to see a list of all tape drives on the system, along with their attribute values. From this list, you can:

- Sort by column headings
- View Alarms (for tape drives displaying a bell-shaped alarm icon) by right-clicking on the tape drive and selecting **Alarms**.
- Perform actions on one or more tape drives (as described in [Recovery Operations Using the OSM Service Connection](#) on page 11-8)

Figure 11-2. OSM: Monitoring Tape Drives Connected to an IOMF2



Note. All tape drives connected to a system appear under the Tape Collection object. When a IOMF2-connected tape drive uses storage routers, those objects appear under that tape drive object in the OSM tree pane hierarchy; however, fibre channel routers appear under the Monitored Service LAN Devices object (after being configured in OSM).

Monitoring Tape Drive Status With SCF

To check the status of all tape drives on your system with SCF:

```
> SCF STATUS TAPE $*
```

A listing similar to this one is sent to your home terminal:

STORAGE - Status TAPE \MINDEN.\$XTAPE				
LDev	State	Primary PID	Backup PID	DeviceStatus
93	STOPPED	1,287	0,279	NOT READY
STORAGE - Status TAPE \MINDEN.\$TAPE0				
LDev	State	Primary PID	Backup PID	DeviceStatus
99	STARTED	1,289	0,278	NOT READY

The data shown in the report means:

LDev	The logical device number
State	The current SCF state of the tape path
SubState	The current SCF substate of the tape path
Primary PID	The primary processor number and process identification number (PIN) of the specified device
Backup PID	The backup processor number and PIN of the specified device
DeviceStatus	The status of the device path

For more information:

- [SCF Object States](#) on page 3-14 describes the possible SCF states of tape drives and other devices.
- The *Guardian User's Guide* provides additional information about tape operations and the tasks you can perform.

Example

To obtain status information about the tape drive \$TAPE0 by using SCF:

```
> SCF STATUS TAPE $TAPE0
```

A listing such as this one is sent to your home terminal:

STORAGE - Status TAPE \MINDEN.\$TAPE0				
LDev	State	Primary PID	Backup PID	DeviceStatus
99	STARTED	1,289	0,278	NOT READY

Monitoring Tape Drive Status With MEDIACOM

The MEDIACOM command STATUS TAPEDRIVE displays the current status of a tape drive. Among other things, this command tells you whether a tape is mounted on the drive, the name of the DEFINE associated with the tape, and which volume catalog and pool owns it.

Note. Manual unloading of a tape is not detected by a tape drive, so information from STATUS TAPEDRIVE can be out of date. For example, STATUS TAPEDRIVE could report that a drive currently has a tape mounted when the tape was removed from the drive by the operator before the command executed.

To check the status of all tape drives on your system with MEDIACOM:

```
> MEDIACOM STATUS TAPEDRIVE
```

A listing similar to this one is sent to your home terminal:

```
MEDIACOM - T6028D42 (18DEC98)

Creating default server.

Tape Drive      Drive  Tape  Tape  Label  Open
Status          Name  Status Type  Mode  Process Name
-----
$XTAPE          DOWN
$TAPE0          FREE
2 tape drives returned.
```

The *DSM/Tape Catalog Operator Interface (MEDIACOM) Manual* explains the fields in this output.

Example

To obtain status information about the tape drive \$TAPE0 by using MEDIACOM:

```
> MEDIACOM STATUS TAPEDRIVE $TAPE0
```

A listing such as this one is sent to your home terminal:

```
MEDIACOM - T6028D42 (18DEC98)

Tape Drive      Drive  Tape  Tape  Label  Open
Status          Name  Status Type  Mode  Process Name
-----
$TAPE0          FREE
1 tape drive returned.
```

Monitoring the Status of Labeled-Tape Operations

Use the `MEDIACOM STATUS TAPEDRIVE` and `STATUS TAPEMOUNT` commands to determine the current status of labeled-tape operations on your system.

For additional information about `MEDIACOM`, the listings it generates, and the tasks it enables you to perform:

- *DSM/Tape Catalog Operator Interface (MEDIACOM) Manual*
- *DSM/Tape Catalog User's Guide*
- *Guardian User's Guide*

Identifying Tape Drive Problems

[Table 11-1](#) lists some of the most common tape drive problems and their possible causes. Additionally, OSM alarm repair actions, degraded attribute values, and EMS event details can help you determine the appropriate cause of action

Table 11-1. Common Tape Drive Problems

Symptom	Problem	Possible Causes
File-system error 48	A security violation has occurred.	An attempted operation was not allowed.
File-system error 49	Various	An unexpired labeled tape was used.
File-system error 51	A runaway tape has been detected.	The system has tried to read a blank tape.
File-system error 66	Various	A hardware failure has occurred, or the tape drive has been purposely brought down.
File-system error 100	A device is not ready.	A tape drive has been brought down, or the drive is not online.
File-system error 195	An operation requires use of <code>\$ZSVR</code> , but it is not running. Tape operation is not allowed.	<code>\$ZSVR</code> has been purposely stopped.
File-system error 218	An interrupt timeout occurs. An I/O process cannot communicate with a tape drive.	A ServerNet addressable controller (SAC) has failed.
No error	A tape label record is missing or incorrect.	An attempt was made to access a tape with a missing or incorrect label.
No error	A tape fails to respond to a <code>BACKUP</code> command.	A tape with an inappropriate label type was mounted in error.
No error	A tape continues to spin beyond the load point.	The load point has fallen off.
No error	Every time a tape is mounted, it is unloaded.	A labeled tape is being mounted in a drive that is open for unlabeled use.

Recovery Operations for Tape Drives

You can perform recovery operations on tape drives using either the SCF interface to the storage subsystem or the OSM Service Connection.

Recovery Operations Using the OSM Service Connection

If the recovery operation calls for an OSM Service Connection action, you can perform an action on one or more tape drive objects.

Performing an OSM Action on a Tape Drive

1. From the OSM Service Connection tree pane (the left-hand pane shown in [Figure 11-1](#) on page 11-3):
 - a. Expand the System and Tape Collection objects to locate the tape drive in need of attention or service.
 - b. Right-click the tape drive object and select **Actions** from the menu.
2. In the Actions dialog box:
 - a. Choose the desired action from the list of available actions.
 - b. Click **Perform action**.
 - c. Check the Action Status window to confirm successful completion of the action, or click **Details** for more information if the action fails.

Also, check the Alarms or Attributes tab to make sure the alarm has been cleared or the degraded attribute value has returned to normal. Using the example in [Figure 11-1](#), you might use the Start action to bring up the selected (highlighted) tape drive. If successful, the Device State should change from Hard Down to Started, and the yellow symbol in both the tree pane and Attributes tab should disappear.

In many cases, there are OSM and SCF equivalents. For example, you can select the OSM Start action or the corresponding SCF START command.

Performing an OSM Action on a Multiple Tape Drives

1. From the Display menu, select **Multi-Resource Actions**.
2. In the Multi-Resource Actions dialog box, select the Tape Drive object to display a list of all tape drives on the system.
3. From this list, select the tape drives upon which you want to perform the action (using the Ctrl key to select multiple tape drives).
4. From the Action drop-down menu, select the desired action.
5. Click **Perform Action**.

Recovery Operations Using SCF

These SCF commands are available for controlling TAPE objects:

SCF Command	Description
PRIMARY	Causes the backup processor of a tape drive to become the primary processor and the primary processor of the drive to become the backup processor.
RESET	Puts a tape drive in a state from which it can be restarted.
START	Initiates the operation of a tape drive.
STATUS	Displays current status information about a tape drive.
STOP	Terminates the operation of a tape drive in a normal manner.

The *SCF Reference Manual for the Storage Subsystem* describes these commands.

Related Reading

For more information about tapes and tape drives, refer to the documentation listed in [Table 11-2](#).

Table 11-2. Related Reading for Tapes and Tape Drives (page 1 of 2)

For Information About...	Refer to...
Tape drives	<i>Integrity NonStop NS-Series Supported Hardware</i>
BACKUP, RESTORE, and BACKCOPY utilities	<i>Guardian Disk and Tape Utilities Reference Manual (for Enscribe and SQL/MP files)</i>
BRCOM utility	<i>Backup and Restore 2.0 Manual (OSS and SQL/MX files)</i>
OSS <code>pax</code> utility	<i>Open System Services Management and Operations Guide (backup and restore of OSS files)</i>
Performing system operations involving tape drives	<i>Guardian User's Guide</i>
Replacing tape drives	Support and Service Library on page 1-12
Recovery operations for generic tape processes	<i>SCF Reference Manual for the Kernel Subsystem</i>
Recovery operations for tape drives	<i>SCF Reference Manual for the Storage Subsystem</i>
Configuring tape drives	<i>SCF Reference Manual for the Storage Subsystem</i>
Starting and stopping tape drives	<i>SCF Reference Manual for the Storage Subsystem</i>
Using the MEDIACOM utility	<i>DSM/Tape Catalog User's Guide</i> <i>DSM/Tape Catalog Operator Interface (MEDIACOM) Manual</i> <i>Guardian User's Guide</i>

Table 11-2. Related Reading for Tapes and Tape Drives (page 2 of 2)

For Information About...	Refer to...
Using the BACKCOPY utility to duplicate backup tapes	<i>Guardian Disk and Tape Utilities Reference Manual</i>
Using the BACKUP utility to save a copy of disk files on tape	<i>Guardian Disk and Tape Utilities Reference Manual</i>
Using the RESTORE utility to copy saved tape files to disk	<i>Guardian Disk and Tape Utilities Reference Manual</i>
Virtual tape server	<i>Virtual Tape Server -- Operations and Administration Guide</i>
Virtual tape server	<i>Virtual Tape Server -- Introduction to Virtual Tape Server</i>
Virtual tape server	<i>Virtual Tape Server -- Installation Guide</i>

12

Printers and Terminals: Monitoring and Recovery

[When to Use This Section](#) on page 12-1

[Overview of Printers and Terminals](#) on page 12-1

[Monitoring Printer and Collector Process Status](#) on page 12-2

[Monitoring Printer Status](#) on page 12-2

[Monitoring Collector Process Status](#) on page 12-2

[Recovery Operations for Printers and Terminals](#) on page 12-3

[Recovery Operations for a Full Collector Process](#) on page 12-3

[Related Reading](#) on page 12-3

When to Use This Section

This section provides a brief overview about monitoring and recovery for printers and terminals. Monitoring printers and terminals, and using the SPOOLCOM utility is discussed more fully in other manuals. Refer to [Related Reading](#) on page 12-3.

Overview of Printers and Terminals

Printers and terminals are attached to the Integrity NonStop server using one of these methods:

- An asynchronous connection provided by the asynchronous wide area network (AWAN) access server for either a terminal or a printer
- An asynchronous connection provided by the ServerNet wide area network (SWAN) concentrator for either a terminal or a printer
- A LAN connection provided by an adapter for a printer

Monitoring Printer and Collector Process Status

This subsection explains how to list the printers on your system and determine their status. It also explains how to check the status of the spooler subsystem collector processes, which accept output from applications and store the output on a disk.

Monitoring Printer Status

To check the status of all printers on your system with the SPOOLCOM utility:

```
> SPOOLCOM DEV
```

A listing similar to this output is sent to your home terminal:

DEVICE	STATE	FLAGS	PROC	FORM
\SAGE.\$S1	WAITING	H	\$SPLX	
\SAGE.\$S2	WAITING	H	\$SPLX	
\AMBER.\$S	WAITING	H	\$SPLP	
\AMBER.\$S2	WAITING	H	\$SPLX	

The value WAITING in the STATE column indicates that the printer is available to print user jobs.

To check the status of the printer \$LASER with the SPOOLCOM DEV command:

```
> SPOOLCOM DEV $LASER
```

A listing such as the following is sent to your home terminal:

DEVICE	STATE	FLAGS	PROC	FORM
\$LASER	WAITING	H	\$SPLP	

The output shows that the printer \$LASER is up and available to print user jobs.

Monitoring Collector Process Status

Check that the collector processes on your spooler subsystem do not become more than about 90 percent full. To check their status:

```
> SPOOLCOM COLLECT
```

A listing similar to this output is sent to your home terminal:

COLLECT	STATE	FLAGS	CPU	PRI	UNIT	DATA FILE	%FULL	
\$S	ACTIVE		0	1	149	4	\$SPOOL.SPL.DATA	40
\$S1	ACTIVE		1	2	149	10	\$SPOOL.SPL.DATA1	28
\$S2	ACTIVE		2	3	149	8	\$SPOOL.SPL.DATA2	0

This listing shows that the three collector processes, \$S, \$S1, and \$S2, are active and none is approaching a full state. The data shown in the report means:

COLLECT	The name of the collector process
STATE	The current state of the collector process, which can be ACTIVE, DORMANT, DRAIN, or ERROR
FLAGS	The current SCF substate of the collector process
CPU	The processor number of the collector process and its backup process
PRI	The execution priority of the collector process (The default value is 145.)
UNIT	The number of 512-word blocks requested by the collector process when it needs more disk space (The default value is 4.)
DATA FILE	The name of the disk file where the collector process stores jobs
%FULL	The percentage of the data file that is full

Recovery Operations for Printers and Terminals

For more information, refer to [Related Reading](#) on page 12-3.

Recovery Operations for a Full Collector Process

If the SPOOLCOM COLLECT display shows any collector process approaching 90 percent capacity, jobs must be deleted from the collector in question.

Related Reading

For more information about printers in your environment, refer to the vendor documentation.

For more information about printers and terminals connected to a SWAN concentrator:

- *WAN Subsystem Configuration and Management Manual*
- *Asynchronous Terminals and Printer Processes Configuration and Management Manual*

For information about the spooler and SPOOLCOM:

- *Guardian User's Guide*
- *Spooler Utilities Reference Manual*

13

Applications: Monitoring and Recovery

[When to Use This Section](#) on page 13-1

[Monitoring TMF](#) on page 13-1

[Monitoring the Status of TMF](#) on page 13-2

[Monitoring Data Volumes](#) on page 13-2

[TMF States](#) on page 13-3

[Monitoring the Status of Pathway](#) on page 13-4

[PATHMON States](#) on page 13-5

[Related Reading](#) on page 13-6

When to Use This Section

This section explains how to monitor the status of the HP NonStop Transaction Management Facility (TMF) and Pathway transaction processing applications. For other applications, such as SQL/MP or SQL/MX, see the appropriate documentation.

Monitoring TMF

This subsection explains how to check the status of TMF and the data volumes it protects. As a system operator, you might check TMF status in your routine system monitoring. You use the TMFCOM command interface to manage and operate TMF.

Monitoring the Status of TMF

To monitor TMF using TMFCOM:

1. At a TACL prompt:

```
> TMFCOM
```

2. At the TMFCOM prompt:

```
~ STATUS TMF
```

Note. The STATUS TMF command presents status information about the audit dump, audit trail, and catalog processes. Thus, in addition to the general TMF information, the STATUS TMF command combines information from the STATUS AUDITDUMP, STATUS AUDITTRAIL, and STATUS BEGINTRANS commands. However, information from the other STATUS commands (STATUS DATAVOLS, STATUS OPERATIONS, STATUS SERVER, and STATUS TRANSACTION) does not appear in the STATUS TMF display.

A TMFCOM report summarizing the current activity of the TMF subsystem, audit trails, and the audit dump and catalog processes is displayed. For example:

```
TMF Status:
System: \SAGE, Time: 6-Jul-1994 11:08:06
State: Started
Transaction Rate: 0.10 TPS
AuditTrail Status:
Master:
  Active audit trail capacity used: 55%
  First pinned file: $MAT1.ZTMFAT.AA000044
  Reason: Active transaction(s).
  Current file: $MAT1.ZTMFAT.AA000045
AuditDump Status:
  Master: State: enabled, Status: active, Process $X545,
  File: $MAT2.ZTMFAT.AA000042
BeginTrans Status: Enabled
Catalog Status:
  Status: Up
```

For an explanation of the TMF state in this display, see [TMF States](#) on page 13-3.

Monitoring Data Volumes

To display information about the data volumes for which the TMF subsystem generates audit records on behalf of transactions performed on those volumes, at a TMFCOM prompt, type:

```
~ STATUS DATAVOLS
```

To control which volumes are displayed, use the STATE, AUDITTRAIL, and RECOVERYMODE parameters.

The normal operating state for a data volume is STARTED, which indicates that the volume is ready to process TMF transactions. Audited requests are allowed for data volumes in this state only where transaction processing is enabled within the subsystem.

For example, to check the status of all data volumes, at a TMFCOM prompt, type:

```
~ STATUS DATAVOLS
```

TMFCOM responds with output similar to:

Volume	Audit Trail	Recovery Mode	State
\$DATA1	MAT	Online	Started
\$DATA2	MAT	Online	Started
\$DATA3	MAT	Online	Recovering
\$DATA4	MAT	Archive	Recovering
\$DATA5	AUX01	Online	Started
\$DATA6	AUX01	Online	Started
\$DATA6	AUX01	Archive	Recovering

TMF States

The TMF subsystem can be in any of the states listed in [Table 13-1](#).

Table 13-1. TMF States (page 1 of 2)

State	Meaning								
Configuring New Audit Trails	The TMF subsystem has not yet been started with this configuration.								
Deleting	The TMF subsystem is purging its current configuration, audit trails, and volume and file recovery information for the database in response to a DELETE TMF command.								
Empty Audit Trail Configuration	The TMF subsystem has been brought up for the first time on this node and thus no configuration exists for it, or a DELETE TMF command was executed.								
Starting	The TMF subsystem is starting and is in one of these conditions: <table border="0" style="margin-left: 20px;"> <tr> <td>Services</td> <td>The subsystem is starting audit-trail service and other services.</td> </tr> <tr> <td>Waiting for Network Transactions to be Resolved</td> <td>The subsystem is waiting for all network transactions to be resolved.</td> </tr> <tr> <td>Data Volumes</td> <td>The TMF subsystem is starting data volumes.</td> </tr> <tr> <td>Running Backout</td> <td>The subsystem is backing out transactions that must be aborted.</td> </tr> </table>	Services	The subsystem is starting audit-trail service and other services.	Waiting for Network Transactions to be Resolved	The subsystem is waiting for all network transactions to be resolved.	Data Volumes	The TMF subsystem is starting data volumes.	Running Backout	The subsystem is backing out transactions that must be aborted.
Services	The subsystem is starting audit-trail service and other services.								
Waiting for Network Transactions to be Resolved	The subsystem is waiting for all network transactions to be resolved.								
Data Volumes	The TMF subsystem is starting data volumes.								
Running Backout	The subsystem is backing out transactions that must be aborted.								
Started	The TMF subsystem has started.								

Table 13-1. TMF States (page 2 of 2)

State	Meaning								
Stopped	The TMF subsystem is stopped.								
Stopping	The TMF subsystem is stopping and is in one of these conditions: <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;">Waiting for Transactions to Finish</td> <td>The subsystem is waiting for all transactions to be finished.</td> </tr> <tr> <td>Data Volumes</td> <td>The subsystem is stopping data volumes.</td> </tr> <tr> <td>Waiting for RDF</td> <td>The subsystem is waiting for the Remote Duplicate Database Facility (RDF) to shut down.</td> </tr> <tr> <td>Services</td> <td>The subsystem is stopping audit-trail service and other services.</td> </tr> </table>	Waiting for Transactions to Finish	The subsystem is waiting for all transactions to be finished.	Data Volumes	The subsystem is stopping data volumes.	Waiting for RDF	The subsystem is waiting for the Remote Duplicate Database Facility (RDF) to shut down.	Services	The subsystem is stopping audit-trail service and other services.
Waiting for Transactions to Finish	The subsystem is waiting for all transactions to be finished.								
Data Volumes	The subsystem is stopping data volumes.								
Waiting for RDF	The subsystem is waiting for the Remote Duplicate Database Facility (RDF) to shut down.								
Services	The subsystem is stopping audit-trail service and other services.								

Monitoring the Status of Pathway

Pathway is a group of related software tools that enables businesses to develop, install, and manage online transaction processing applications. Several Pathway environments can exist for a system. As a system operator, you might check the status of Pathway in your routine system monitoring. This subsection explains how to check the status of the Pathway transaction processing applications.

1. To determine the names of the Pathway processes running on your system:
 - > STATUS *, PROG \$*.*.PATHMON
2. To access PATHCOM to communicate with one of the PATHMON processes:
 - > PATHCOM *\$pathmon-process-name*
3. At the PATHCOM prompt:

```
= STATUS PATHWAY
```

For example, to check the status of the PATHMON process for the Pathway environment on your system:

```
> PATHCOM $ZVPT
```

```
$Y290: PATHCOM - T9153D20 - (01JUN93)
COPYRIGHT TANDEM COMPUTERS INCORPORATED 1980 - 1985, 1987 -
1992
```

```
= STATUS PATHWAY
```

PATHCOM responds with output such as:

	RUNNING					
EXTERNALTCPS	0					
LINKMONS	0					
PATHCOMS	1					
SPI	1					
	RUNNING	STOPPED	THAWED	FROZEN	FREEZE	PENDING
SERVERCLASSES	13	5	18	0		0
	RUNNING	STOPPED	PENDING			
SERVERPROCESSES	13	40	0			
TCPS	1	0	0			
	RUNNING	STOPPED	PENDING	SUSPENDED		
TERMS	1	0	0	0		

This output provides information about the number of Pathway processes and servers that are running, stopped, and so forth.

- To check the state of the PATHMON process within the Pathway environment and its status for your application:

```
= STATUS PATHMON
```

PATHCOM responds with output such as:

PATHMON -- STATE=RUNNING CPUS 6:1					
PATHCTL (OPEN) \$GROG.VIEWPT.PATHCTL					
LOG1 SE (OPEN) \$0					
LOG2 (CLOSED)					
REQNUM	FILE	PID	PAID	WAIT	
1	PATHCOM	\$Y622	8,001		
2	TCP	\$Y898			

PATHMON States

The status of the PATHMON process can be either STARTING or RUNNING:

- STARTING indicates that a system load or cool start has not finished.
- RUNNING indicates that a system load or cool start has finished.

The other elements of the STATUS PATHMON output are:

- CPUS shows the number of the primary and backup processors in which the PATHMON process is running. If the backup PATHMON process is not running, the second number is blank.
- PATHCTL, LOG1, and LOG2 contain information about the PATHMON control file and the logging files.
- The REQNUM column contains the PATHMON internal identifiers of application requesters that are currently running in this environment.
- The FILE column identifies the type of requester.

- The WAIT column indicates whether the process is waiting, which can be caused by one of these conditions:
 - IO The request is waiting for an I/O operation to finish.
 - LOCK The request is waiting for an object that has been locked by another requester.
 - PROG-DONE The request is waiting for a RUN PROGRAM to finish.

Related Reading

For more information about Pathway or interpreting displays, refer to:

- *TS/MP System Management Manual*
- *TMF Operations and Recovery Guide*
- *TMF Planning and Configuration Guide*
- *TMF Reference Manual*

14

Power Failures: Preparation and Recovery

[When to Use This Section](#) on page 14-2

[System Response to Power Failures](#) on page 14-2

[NonStop NS-Series Cabinets \(Modular Cabinets\)](#) on page 14-2

[External Devices](#) on page 14-2

[ESS Cabinets](#) on page 14-3

[Air Conditioning](#) on page 14-3

[Preparing for Power Failure](#) on page 14-3

[Set Ride-Through Time](#) on page 14-3

[Monitor Power Supplies](#) on page 14-4

[Monitor Batteries](#) on page 14-4

[Maintain Batteries](#) on page 14-4

[Power Failure Recovery](#) on page 14-4

[Procedure to Recover From a Power Failure](#) on page 14-5

[Setting System Time](#) on page 14-5

[Related Reading](#) on page 14-5

When to Use This Section

Use this section for information about how to prepare for power failures and how to recover if a power failure occurs.

System Response to Power Failures

NonStop NS-Series Cabinets (Modular Cabinets)

NonStop NS-series servers are designed to operate in a computer-room environment containing a site UPS. Without a UPS, a system will stop uncontrollably when the power is lost. An optional UPS module can be installed in a modular cabinet to provide power if no site UPS is available.

If AC power to a NonStop NS-series server is lost, the system will shut down after a preset time as long as UPS power is available. If the system shuts down, it will be necessary to restart the system manually when AC power is restored.

If modular cabinets lose power without a UPS, recovery of lost data will be difficult and files might be corrupted.

Regardless of system power-fail scenario, if site air conditioning fails and the computer-room temperature rises, the system might shut itself down uncontrollably as each processor reaches its critical temperature. Refer to the *NonStop NS-Series Site Preparation Guide*.

NonStop S-Series I/O Enclosures

NonStop S-series I/O enclosures have internal batteries and do not require a UPS. NonStop S-series enclosures must shut down before their battery power is lost. For information about power fail for I/O enclosures, refer to the *NonStop NS-Series Planning Guide*.

External Devices

External (peripheral), devices, such as tape drives, external disk drives, LAN routers, and SWAN concentrators, are not backed up by internal batteries. External devices behave differently than a system during power failure according to the UPS support provided:

No UPS	Peripheral devices will fail immediately if not supported by a UPS.
With a Site UPS	A site UPS should support external devices until its capacity to supply power is exhausted.
With an Internal UPS	An optional internal UPS will not support external devices during power failure.

During a power failure, a ServerNet/DA remains operational during the power-fail delay time, but the external modular disk and tape subsystems attached to it do not. This type of situation could result in data-integrity problems if the system software continues processing data from an external disk drive or tape drive during a short power outage.

If a power failure occurs and the processors resume operations but one or more external devices fail, data integrity problems can occur. The application programs must be resilient to such device failures.

ESS Cabinets

ESS cabinets require a site UPS. ESS cabinets are not powered off programatically with a power failure because they may be attached to other systems.

Air Conditioning

Unless the site UPS provides for maintaining air conditioning, it is possible that the temperature could rise in the computer room complicating the ability of a system to survive overheating before shutting itself down. For information about site overheating problems, refer to the *NonStop NS Series Site Preparation Guide*.

Preparing for Power Failure

To prepare for power failures, set ride-through time, configure OSM power fail support, and regularly monitor power supplies and batteries.

Set Ride-Through Time

Ensure that the system is set for the proper ride-through time. The default powerfail delay time for NS-series systems that are configured with rack-mounted UPSs is 30 seconds. Contact HP Expert Services for the optimum ride-through time for your system.

The ride-through time is set using the SCF commands to alter the `POWERFAIL_DELAY_TIME` parameter. Refer to the *SCF Reference Manual for the Kernel Subsystem*.

Configure OSM Power Fail Support

You must also configure OSM power fail support by performing a **Configure Power Source as UPS** action for at least one processor switch power supply unit (for Integrity NonStop NS16000 systems) or one IOAM or VIO power supply unit (for Integrity NonStop NS14000 or NS1000 systems). Configuring at least one power supply as being powered by UPS in the event of a power outage causes OSM to monitor a power outage and, if the AC power is not restored before the specified ride-through time period expires, OSM initiates an orderly system shutdown.

For more information, see the *NonStop NSxxxx Hardware Installation Manual* for your Integrity NonStop NS16000, NS14000, or NS1000 server or the *OSM Service Connection User's Guide*.

Monitor Power Supplies

Monitor power-generating equipment and run regular checks on any backup generators to make sure that you can handle extended power outages.

Monitor Batteries

- Monitoring site UPS batteries is the responsibility of the customer. OSM does not interface with a site UPS or batteries.
- Monitoring batteries in an internal UPS and ERM is performed using OSM.
- Monitor batteries in I/O enclosures using OSM. For more information on battery attributes and actions, see the OSM Service Connection online help.

Maintain Batteries

Make sure that all installed batteries and spare batteries are always fully charged. Correct any problems that are causing the batteries to become drained.

In general, batteries are constantly charging when AC power is available to a system. When recovering from a power failure event, it will take time for the batteries to be fully charged. For different kinds of batteries:

- Maintenance and charging of site UPS batteries is the responsibility of the customer.
- Information about maintenance and charging of batteries in an optional internal UPS and ERM is located in documentation that comes with the products.
- Information about maintenance and charging of batteries in I/O enclosures is located in the [Support and Service Library](#) for NonStop S-series servers. Spare batteries for I/O enclosures require charging every six months.

Power Failure Recovery

After a power failure, if AC power is restored to a NonStop NS-series server while the batteries are still holding up the system, it will not be necessary to restart the system. Resume operations.

Depending on the configuration of UPS resources, power failure can last long enough to leave the system with some processors down because the batteries were drained to the point where the processors can no longer operate. In conjunction with an internal UPS, a homogeneous power off can be configured to avoid this condition.

If the power failure lasts long enough to drain the batteries completely, the system Power/Fail should stop the system. When power is restored, the operator must then restart the system.

Procedure to Recover From a Power Failure

After power is restored:

1. Power on the system using OSM LLL described in [Powering On a System](#) on page 15-2.
2. Log on to the OSM Service Connection and check the status of all system components to make sure they are started.
3. Use SCF commands to check the status of external devices and, if necessary, to restart any external devices to bring them back online.

Setting System Time

Setting the system time is not normally required following a power failure. System time is maintained by a time-of-day battery in the p-switch, IOAM, or VIO logic board that is not affected by a power outage. If required, however, you can set the system time, either programmatically or by using the TACL command interpreter. Refer to the *Guardian Procedure Calls Reference Manual* or the *TACL Reference Manual*.

Related Reading

For more information about preparing for and recovering from power failures:

- The effect of power failures on NonStop NS-series servers, see the *NonStop NS-Series Planning Guide*.
- The ride-through time, see the *SCF Reference Manual for the Kernel Subsystem*.
- The TACL SETTIME command, see the *TACL Reference Manual*.
- Setting system time programmatically, see the *Guardian Procedure Calls Reference Manual*.
- Removing, installing, and recycling batteries, see the documentation provided for the type of batteries used in the UPSs.

15

Starting and Stopping the System

[When to Use This Section](#) on page 15-2

[Powering On a System](#) on page 15-2

[Powering On the System From a Low Power State](#) on page 15-3

[Powering On the System From a No Power State](#) on page 15-3

[Starting a System](#) on page 15-5

[Loading the System](#) on page 15-5

[Starting Other System Components](#) on page 15-9

[Performing a System Load](#) on page 15-9

[Performing a System Load From a Specific Processor](#) on page 15-11

[Reloading Processors](#) on page 15-12

[Minimizing the Frequency of Planned Outages](#) on page 15-14

[Anticipating and Planning for Change](#) on page 15-14

[Stopping Application, Devices, and Processes](#) on page 15-14

[Stopping the System](#) on page 15-16

[Alerts](#) on page 15-16

[Halting All Processors Using OSM](#) on page 15-16

[Powering Off a System](#) on page 15-17

[System Power-Off Using OSM](#) on page 15-17

[System Power-Off Using SCF](#) on page 15-17

[Emergency Power-Off Procedure](#) on page 15-18

[Troubleshooting and Recovery Operations](#) on page 15-18

[Fans Are Not Turning](#) on page 15-18

[System Does Not Appear to Be Powered On](#) on page 15-19

[Green LED Is Not Lit After POSTs Finish](#) on page 15-19

[Amber LED on a Component Remains Lit After the POST Finishes](#) on page 15-19

[Components Fail When Testing the Power](#) on page 15-19

[Recovering From a System Load Failure](#) on page 15-20

[Getting a Corrupt System Configuration File Analyzed](#) on page 15-21

[Recovering From a Reload Failure](#) on page 15-21

[Exiting the OSM Low-Level Link](#) on page 15-22

[Opening Startup Event Stream and Startup TACL Windows](#) on page 15-22

[Related Reading](#) on page 15-24

When to Use This Section

Normally, you leave a system running. However, some procedures or recovery actions require you to start the system (perform a system load) or stop or power off the system.

- Stop and then power off a system before:
 - An extended planned power outage for your building or computer room.
 - Performing some major maintenance or repair operations as noted in the documentation.
- Stop or restart a system (without powering off) when:
 - Installing an updated RVU or some software product revisions (SPRs).
 - Performing some recovery operations as noted in the documentation.
 - Restarting the system after the entire system has been shut down with the operating system images and files on disk still intact.

Powering On a System

Powering on a system delivers AC power to the system cabinets. Fans on the processor switches, processor Blade Elements, IOAM or VIO enclosures, and disk drive enclosures start turning, and air begins to circulate through the components. After the fans start to operate, most other system components begin to power on.

Status light-emitting diodes (LEDs) on the system components light during a series of power-on self-tests (POSTs). Any of the LEDs can become lit briefly during the POSTs. After all the POSTs finish successfully, which might take up to 10 minutes, only the green power-on LED on each component in the system enclosures should remain lit. For more information about the LEDs, refer to [Using the Status LEDs to Monitor the System](#) on page 3-20.

The method you use to power on the system depends on whether the system is in a low power state or completely powered off:

- [Powering On the System From a Low Power State](#) on page 15-3
- [Powering On the System From a No Power State](#) on page 15-3

Powering On the System From a Low Power State

1. Log on to the OSM Low-Level Link.
2. From the tree pane, right click the system and select actions.
3. Select **Power On System**.
4. If your maintenance LAN is not configured with the dynamic name service (DNS) or does not have reverse look-up, you must perform a hard reset of the maintenance entities (MEs) in each p-switch or IOAM enclosure, or the integrated maintenance entities (IMEs) in each VIO enclosure*:
 - a. From the Log On to HP OSM Low-Level Link dialog box, select Logon with Host Name or IP Address.
 - b. Enter the IP address of:
 - a maintenance entity (ME) in a p-switch
 - an ME in an IOAM enclosure
 - an IME in a VIO enclosure
 - c. Expand the tree pane to locate the ME or IME.
 - d. Right-click that ME or IME object and select **Actions**.
 - e. Select **Hard Reset**.
 - f. Click **Perform Action**.
 - g. A message appears:

```
Hard Reset action will make the current session lost.
After OSM Low-Level Link completes Hard Reset action, it
will log you off. Do you really want to reset the ME?
```
 - h. Click OK.
 - i. Close the Low-Level Link.
 - j. Repeat these steps for the other ME(s) or IME.

* For Integrity NonStop NS14000 and NS1000 systems, there are no p-switches. Instead, there are either two MEs in the IOAM enclosure or an IME in each of the two VIO enclosures.

Powering On the System From a No Power State

To power on the system when you first receive your system, you refer to the *NonStop NSxxxx Hardware Installation Manual* for your Integrity NonStop NS16000, NS14000, or NS1000 server. To power on the system subsequently:

1. Before you power on any system enclosures, power on the external system devices and any other devices you want started when the system starts. External system devices include tape devices, Enterprise Storage Systems (ESSs), printers, and terminals.

Refer to the documentation that accompanies the device for instructions on powering on. For example:

- You must power on Fibre Channel to SCSI Converter devices connected to your system before you power on the tape devices attached to it. The converter must be powered on first to be able to discover the tape devices as they are powered on.
 - Maintenance switches installed outside of a modular cabinet must be powered on according to the instructions provided with the switch.
2. Locate the circuit breaker that controls the power cords connected to the system cabinets. In most case, each modular cabinet in a system powers on as soon as soon as the power is applied. In addition, these components power on when the Integrity NonStop NS-server is powered on:
 - NonStop I/O enclosures connected to the system
 - Maintenance switches installed in a modular cabinet power
 3. To physically monitor power-on activity:
 - a. Check fan activity for the processor switches, processor Blade Elements, IOAM or VIO enclosures, and disk drive enclosures. Check that the fans are turning and that you can feel air circulate through the components.
 - b. After the POSTs finish, check that only green power-on LEDs are lit in the system components before you start the server. For more information about status LEDs, see [Using the Status LEDs to Monitor the System](#) on page 3-20.
 4. After the POSTs finish, check the AC power cords. Perform this test only if you have connected redundant power cords to separate circuits:
 - a. If you have a UPS installed, switch off the UPS outputs.
 - b. If you do not have a UPS installed:
 1. Locate the circuit breaker that controls half the power cords.
 2. Switch this breaker off.
 - c. Check that all components are still operating.

Note. The maintenance switch does not have redundant power.

 - d. Switch this breaker back on.
 - e. Locate the other circuit breaker that controls the other half of the power cords.
 - f. Switch this breaker off.
 - g. Check that all components are still operating.
 - h. Switch this breaker back on.

- i. If any components fail during any of the power shutdowns, see [Components Fail When Testing the Power](#) on page 15-19.
 - j. If you have a UPS installed, switch off the UPS outputs.
5. If you have a UPS installed, check that the UPS is fully charged. Then, test the UPS by turning off both circuit breakers.
6. Log on to the OSM Low-Level Link.
7. Select **System Discovery**.
8. Double-click the System.
9. Double-click each Group 40*n*; for example, Group 400.
10. Check that the logical processors are displayed.
11. Double-click each Group 1*nn*; for example, Group 100.
12. For each Group 1*nn*, check that module 2 and module 3 are displayed.
13. If any of these components are not yet displayed, wait before you start the system.
14. After the system is powered on, you must wait 5 minutes before starting the system.
15. You can now start your system as described in [Starting a System](#) on page 15-5.

Starting a System

Use the OSM Low-Level Link to start a system. Starting a system includes a system load from disk into the memory of one processor; followed by the reload of the remaining processors.

Loading the System

Perform a system load to load the NonStop operating system.

Alerts

- All processors in the system must be in a halted state before you perform a system load.
- To perform processor dumps during a system load, see the considerations in [System Load to a Specific Processor](#) on page 15-6.

Normal System Load

Normally, you initiate the system load as described in [Performing a System Load](#) on page 15-9. When you choose this method:

- Processor 0 or 1 is loaded. See [System Load Paths for a Normal System Load](#) on page 15-7.
- The remaining processors are primed for reload.
- Two startup event stream windows and two startup TACL windows are automatically launched on the system console configured to receive them.
- The CIIN function is enabled by default. See [CIIN File](#) on page 16-5.

System Load to a Specific Processor

Alternately, you can perform a system load from a specified processor. When you load the system from a specified logical processor:

- The other logical processors are not primed automatically.
- Because the processors not being loaded are not primed, you can perform processor dumps, if necessary. If you need to dump processors, refer to [Section 9, Processors and Components: Monitoring and Recovery](#).
 - Do not prime or reset all the processor elements in a logical processor until after the memory dump.
 - Disable the CIIN file to prevent any processors configured to reload in that file from being reloaded. The startup event stream windows and Startup TACL windows are not launched automatically when the CIIN file is disabled. See [CIIN File](#) on page 16-5.
- The Processor Element Dump Setting option becomes available in the System Load dialog box. For more information about this option, see [Section 9, Processors and Components: Monitoring and Recovery](#).
- You must initiate the load action as described in [Performing a System Load From a Specific Processor](#) on page 15-11.

System Load Disks

An Integrity NonStop NS-series system can contain multiple system disk pairs in different locations. Use the System Load dialog box to select which system disk to load from. You select the system disk from the Configuration drop-down menu. The system load disk you choose must be in the configured location, properly configured as a system disk, and contain the software configuration that you want to load. The system disk that you load from starts as \$SYSTEM. Any alternate system disks start using their alternate name.

After you select a system load disk, the Disk Type box indicates whether you've selected a Fibre Channel (FCDM) or SCSI disk. The Path window is populated with information about four load paths. You can double-click on a row to make changes. However, the changes do not persist after the dialog box is closed.

You can choose these system load disks:

- An FCDM-Load attempts to load the system from a system disk in the disk drive enclosure connected to IOAM enclosure group 110:

Path	IOAM		FCSA		Disk Drive Enclosure	
	Group	Module	Slot	SAC	Shelf	Bay
Primary	110	2	1	1	1	1
Backup	110	3	1	1	1	1
Mirror	110	3	1	2	1	1
Mirror Backup	110	2	1	2	1	1

Note. For Integrity NonStop NS14000 and NS1000 servers, Fibre Channel disks are connected to IOAMs or VIO enclosures located in group 100. For more information, see the *NonStop NSxxxx Hardware Installation Manual* for your Integrity NonStop NS14000 or NS1000 server, or the *Versatile I/O (VIO) Manual*.

- A SCSI-Load attempts to load the system from a disk in group, module, slot 11.1.11 of a NonStop S-series I/O enclosure.
- A load from \$SYSTEM attempts to load the system from a disk in group, module, slot 11.1.11 of a NonStop S-series I/O enclosure, by default.
- You can configure additional alternate system disks to load from. To create an alternate system disk, see the *NonStop NSxxxx Hardware Installation Manual* for your Integrity NonStop NS16000, NS14000, or NS1000 server. Then use OSM to make the disk available in the Configuration Drop-down menu in the System Load dialog box.

System Load Paths for a Normal System Load

16 paths are available for loading. [Table 15-1](#) describes each load path in order of use. The system load task attempts to use each path until the system load is successful or all possible paths have been tried. If the system load fails along all paths, refer to [Troubleshooting and Recovery Operations](#) on page 15-18.

Table 15-1. System Load Paths in Order of Use

Load Path	Description	Data Travels		
		From	To Processor	Over ServerNet Fabric
1	Primary	\$SYSTEM-P	0	X
2	Primary	\$SYSTEM-P	0	Y
3	Backup	\$SYSTEM-P	0	X
4	Backup	\$SYSTEM-P	0	Y
5	Mirror	\$SYSTEM-M	0	X

Table 15-1. System Load Paths in Order of Use

Data Travels				
6	Mirror	\$\$SYSTEM-M	0	Y
7	Mirror backup	\$\$SYSTEM-M	0	X
8	Mirror backup	\$\$SYSTEM-M	0	Y
9	Primary	\$\$SYSTEM-P	1	X
10	Primary	\$\$SYSTEM-P	1	Y
11	Backup	\$\$SYSTEM-P	1	X
12	Backup	\$\$SYSTEM-P	1	Y
13	Mirror	\$\$SYSTEM-M	1	X
14	Mirror	\$\$SYSTEM-M	1	Y
15	Mirror backup	\$\$SYSTEM-M	1	X
16	Mirror backup	\$\$SYSTEM-M	1	Y

Configuration File

Normally, you select Current (CONFIG), the default system configuration file. For the system disk volume you select to load the system, CONFIG represents the system configuration database that is currently running or was last running.

If you cannot load the system using Current (CONFIG), you might need to use these files to recover:

- **Saved Version (CONF_{xxyy})** is a saved system configuration file. Use this file to recover from a configuration change that causes a problem. If you cannot load the system using the CONFIG file, you can use a saved version in the form *xx.yy*, if one is available. For example, the CONF0205 file would be specified as 02.05. Use this method to recover from a configuration change that caused a problem (such as a system freeze). When the system starts and displays a TACL prompt, you can log on and start the rest of the system applications.
- **Base (CONFBASE)** is the most basic configuration required for system startup. Although you will probably never need to load the system from the CONFBASE file, you might need to use this file if you cannot load the system using any other method.

For more information about when to use CONFSAVE or CONFBASE, see [Recovering From a System Load Failure](#) on page 15-20

Starting Other System Components

HP recommends that you bring your system up in stages, verifying each stage, to facilitate recovery if any step fails. When the system starts, many individual devices, processes, applications, and communications lines start automatically, but others might need to be started using start up files. Follow your site's procedures for starting your applications.

- Many processes are configured by default to be started automatically by the \$ZPM persistence monitor. These processes include the Kernel subsystem, SLSA subsystem, storage subsystem, and WAN subsystem. The manager processes for these subsystem start disks, SWAN concentrators, some LAN devices, and many processes.
- You can use SCF to configure other processes, typically monitor or manager processes, to start automatically as generic processes when the system starts. For example, you can use the SCF interface to the Kernel subsystem to add these processes to the system configuration database:
 - \$ZEXP, the Expand manager process
 - \$ZPMON, the OSS monitor process

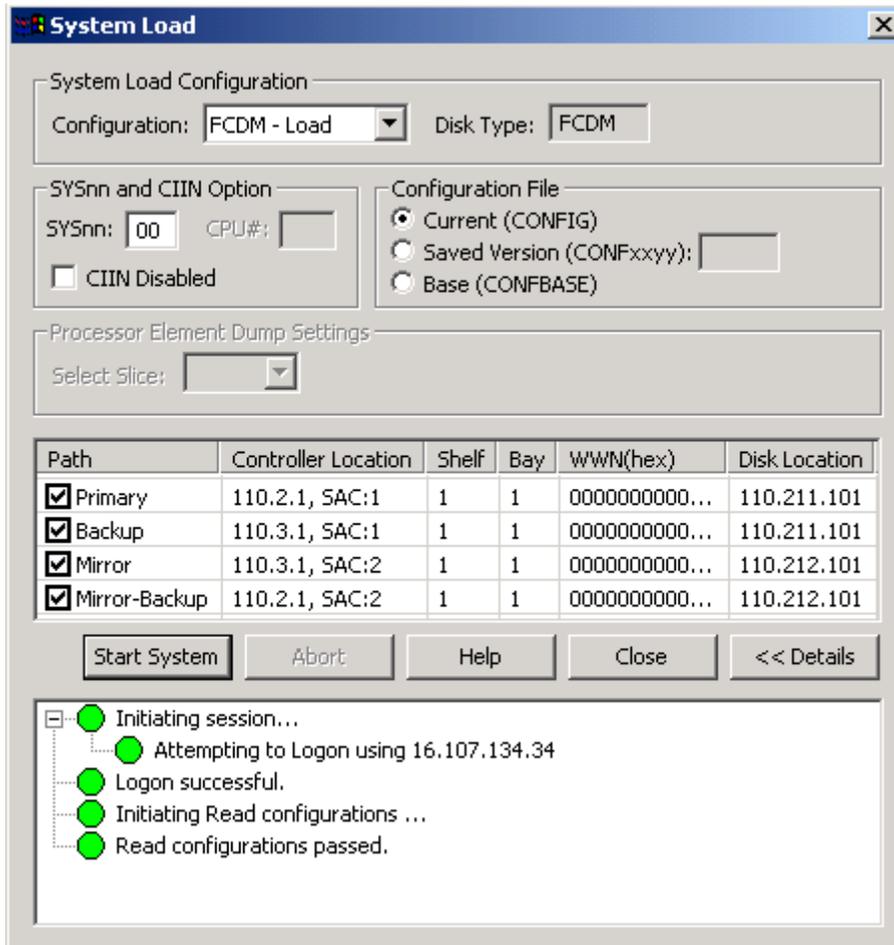
For more information about configuring generic processes to start automatically, refer to the documentation in [Related Reading](#) on page 15-24.

- You can include commands in startup command files that you invoke from a TACL prompt or another startup file. For some techniques to make startup command files run as efficiently as possible, refer to [Writing Efficient Startup and Shutdown Command Files](#) on page 16-9.

Performing a System Load

To perform a normal system load:

1. Verify that all processors are halted as described in [Stopping the System](#) on page 15-16. All processors in the system must be halted before you initiate a system load.
2. Log on to the OSM Low-Level Link.
3. From the OSM Low-Level Link toolbar, click **Start system**.
4. If you initiate a system load while processors are running, a message appears asking whether you want to proceed. If you click Yes, all the processors are halted, then the system load begins.

Figure 15-1. System Load Dialog Box

5. From the Configuration drop-down menu under System Load Configuration, select a system load volume. You can select the \$SYSTEM, FCDM-Load, SCSI-Load, or an alternate system volume.
6. In the SYS nn field, enter the number of the SYS nn subvolume. The value nn must be a two-digit octal number in the range 00 through 77.
7. In the Configuration File box, select a system configuration file. In most cases, you should select the Current (CONFIG) file.
8. Select or clear the CIIN disabled check box. For a normal system load, check that the CIIN disabled check box is cleared so that the commands in the CIIN file execute.
9. To make changes to the load paths, double-click on a row in the Path window.
10. Click **Start system**.

11. Check for messages in these windows and dialog boxes:
 - System Load dialog box
 - The primary startup event stream window (startup event stream messages) and the primary startup TACL window (startup messages)
12. After the System Load dialog box displays the “System Startup Complete” message, close the dialog box.
13. From the Processor Status dialog box, check the status of all processors. At least one processor must be running. Determine whether you need to reload any remaining processors. If needed, reload any remaining processors. See [Reloading Processors](#) on page 15-12.

Performing a System Load From a Specific Processor

Use this method if you need to dump processor memory. See [Section 9, Processors and Components: Monitoring and Recovery](#). To perform a system load into a specified processor:

1. From the OSM Low-Level Link toolbar, click **Processor Status**. The Processor Status dialog box appears.
2. In the Processor Status dialog box:
 - a. Select the processor you want to load.
 - b. From the Processor Actions drop-down menu, select **Load**.
 - c. Click **Perform action**. The System Load dialog box appears.
3. In the System Load dialog box, from the Configuration drop-down menu under System Load Configuration, select a system load volume. You can select the current \$SYSTEM, FCDM-Load, SCSI-Load, or an alternate system volume.

After you select a system load disk, the Disk Type box indicates whether you've selected a Fibre Channel (FCDM) or SCSI disk. The Path window is populated with information about four load paths. You can double-click a row to make changes for this system load. The changes do not persist after the dialog box is closed.
4. In the SYS nn field, enter the number of the SYS nn subvolume. The value nn must be a two-digit octal number in the range 00 through 77.
5. In the Configuration File box, select a system configuration file. Normally, you choose Current (CONFIG). See [Configuration File](#) on page 15-8.
6. Select or clear the CIIN disabled check box. If you need to dump processors, you must select the CIIN disabled option.
7. If you need to perform a memory dump of the logical processor you are loading, you must omit a Blade Element to prevent the processor element selected from being loaded or primed. From the Select Blade Element drop-down menu in the Processor Element Dump Settings field, select the Blade Element.

8. Click **Load**.
9. Check for messages in the System Load dialog box. After the “System Startup Complete” message, close the dialog box.
10. In the Processor Status dialog box, check the status of all processors. At least one processor must be running. Determine whether you need to reload any remaining processors.
11. Dump processor memory, if needed. For more information about dumping processor memory, refer to [Section 9, Processors and Components: Monitoring and Recovery](#).
12. If needed, reload any remaining processors.

Reloading Processors

After the system load, use one of these procedures to reload the remaining processors in the system:

- [Reloading Processors Using the RELOAD Command](#) on page 15-12
- [Reloading Processors Using OSM](#) on page 15-12

To reload a halted processor and perform memory dumps, use the reload procedures in [Section 9, Processors and Components: Monitoring and Recovery](#).

Reloading Processors Using the RELOAD Command

1. From a TACL prompt, log on to the system as a super user ID (255,255).
2. Reload the remaining processors. For example:

```
> RELOAD (01 - 15), PRIME
```
3. Check that the reload initiated successfully. This message appears in the TACL window:

```
PROCESSOR RELOAD: nn
```

For more information about using the RELOAD command, refer to the *TACL Reference Manual*.

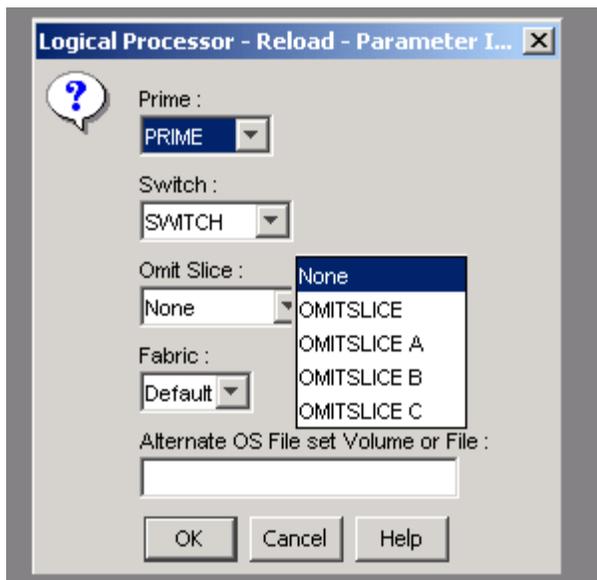
Reloading Processors Using OSM

The OSM Service Connection provides a Reload action on the Logical Processor object. You can perform the action on a single or multiple processors. The OSM action lets you reload an entire processor or omit a Blade Element from the reload action so you can dump the PE for that Blade Element before reintegrating it into the running processor. To reload a single processor, see [Section 9, Processors and Components: Monitoring and Recovery](#).

To reload a multiple processors, use the Multi-Resource Actions dialog box (available from the Display menu of the OSM Service Connection):

1. In the Multi-Resource Actions dialog box, select **Logical Processor** from the Resource Types list. All Logical Processors in the system will be displayed in the right-hand pane.
2. Select **Reload** from the Action list under Selection Criteria.
3. From the list of Logical Processors, select all the processors and click **Add** to move them to the lower list (you can select and add one at a time, or Ctrl-click to select more than one).
4. Once all the processors you want to reload are in the lower list, click **Perform Action**.
5. Click **OK** to dismiss the confirmation dialog box.
6. In the Logical Processor Reload Parameters dialog box (See [Figure 15-2](#) on page 15-13), select the appropriate options. See OSM online help for information about the options.
7. Click **OK**.
8. Check that the reload initiated successfully. From the Low-Level Link, the Processor Status dialog box shows the status for the processor as “Executing NonStop OS.”

Figure 15-2. Logical Processor Reload Parameters



Minimizing the Frequency of Planned Outages

To minimize the frequency of planned outages:

- Anticipate and plan for change
- Perform changes online

Anticipating and Planning for Change

Anticipating and planning for change is a key requirement for maintaining an enterprise-level, 24 x 7 operation. To avoid taking a NonStop NS-series system down unnecessarily:

- Evaluate system performance and growth—Track system usage and anticipate system capacity and performance requirements as new applications are introduced.
- Provide adequate computer room resources—Avoid unnecessary downtime by ensuring you have enough physical space and power and cooling capacity to handle future growth.
- Configure the system with change in mind—Configure the system in a way that easily accommodates future growth. One way to do this is to select limits that allow for growth. For example, by configuring enough objects to provide for the anticipated growth of your online transaction processing environment, you can increase the maximum number of objects controlled by PATHMON objects without a system shutdown.

Stopping Application, Devices, and Processes

Whenever possible, schedule system shutdowns in advance so that system users are prepared. Then, stop applications, devices, and processes in an orderly fashion. To include shutdown commands in a shutdown file, see [Stopping the System](#) on page 15-16.

Unless you stop a system in a careful and systematic manner, you can introduce abnormalities in the system state. Such abnormalities can affect disk file directories and can cause the processors to hang in an endless loop when you attempt to load your system.

You must be aware of which processes must not or cannot be stopped. For example, some TCP processes must not be stopped. System processes must not be stopped. Generic processes configured to be persistent cannot be stopped.

Note the effect on the system when you stop these applications:

- Stopping Pathway applications begins shutdown of all TCP objects (shutting down TERM objects and then themselves) in parallel. New work is disallowed. The PATHMON process logs the start and completion of SHUTDOWN2. It does not log status messages during shutdown.

- Stopping DSM/SCM stops the CNFGINFO server process, the Pathway environment for DSM/SCM, the alternate EMS collector \$ZPHI, and TCP/IP processes for DSM/SCM, as this example shows:

```
STOP CNFGINFO server process $ZPHC
STOP DSM/SCM Pathway system $YPHI
PATHCOM $YPHI;SHUTDOWN !,WAIT
$Z02H: TCP TCP-H, STOPPED
$Z02H: TCP TCP-T, STOPPED
STOP DSM/SCM Alternate EMS Collector $ZPHI
```

- Following the SPOOLER DRAIN command, the collectors allow current jobs to finish but reject new opens with a file-system error 66 (device downed). When you drain the spooler, each collector stops when it has no more open jobs. Each print process finishes printing any active jobs and then stops. After all collectors and print processes have stopped, the supervisor stops. The spooler enters the dormant state, ready to be warm started.
- Following the SCF CONTROL, DISK REFRESH command, all other disk I/O is suspended. The amount of time a refresh operation takes to finish depends on the amount of disk cache containing dirty pages in use at the time, and writing to disk can take several minutes.

Stop processes and applications in this order:

1. After you send a message alerting users of the shutdown, stop all user applications.
2. If your system is equipped with Pathway, stop Pathway applications. At the Pathway prompt:


```
= SHUTDOWN2, MODE ORDERLY
```
3. Stop Distributed Systems Management/Software Configuration Manager (DSM/SCM) if it is running. At a TACL prompt:
 - a. Type this VOLUME command:


```
> VOLUME $DSMSCM.ZDSMSCM
```
 - b. Stop DSM/SCM:


```
> RUN STOPSCM
```
4. Stop communications lines, such as Expand lines.
5. Identify and stop any remaining processes that should be stopped individually:
 - a. Use the TACL PPD and STATUS commands to help you identify running processes.
 - b. Use the TACL STOP command to stop running processes.

6. Drain the spooler. The SPOOLCOM SPOOLER DRAIN command stops the spooler in an orderly manner. It is the only recommended way to stop the spooler. At a TACL prompt:

```
> SPOOLCOM supervisor-name, SPOOLER, DRAIN
```

7. Stop the TMF subsystem. At the TACL prompt:

```
> TMFCOM STOP TMF
```

8. Refresh the disks to put them in an orderly state before shutdown. Use the SCF CONTROL DISK, REFRESH command:

```
> SCF CONTROL DISK $*,REFRESH
```

9. If the system is a member of a ServerNet cluster, HP recommends that you first remove the system from the cluster. To remove the system from the cluster, refer to the *ServerNet Cluster 6780 Operations Guide* (for 6780 switches) or the *ServerNet Cluster Manual* (for 6770 switches).

Stopping the System

Stopping a system halts each processor (terminating all processes running in each processor) on the system in an orderly fashion.

Alerts

Before stopping a system:

- Stop applications, devices, and processes in an orderly fashion. See [Stopping the System](#) on page 15-16.
- To maximize application availability, make stopping the system a planned event whenever possible.

Halting All Processors Using OSM

To place processors in a halt state and set the status and registers of the processors to an initial state:

1. Prepare the system by shutting down applications and performing all tasks in [Stopping the System](#) on page 15-16.
2. Log on to the OSM Low-Level Link.
3. On the toolbar, click **Processor Status**.
4. In the Processor Status dialog box, select all processors to be halted.

To select multiple processors, use the Shift key, but the processors must be in numerical order. For example, you can select processors 2, 3, and 4, but not 2 and 4.

5. From the **Processors Actions** menu, select **Halt**.
6. Click **Perform action**.
7. A message box asks whether you are sure you want to perform a halt on the selected processors. Click **OK**.

Powering Off a System

Use OMS or SCF to place most system components in a low-power state before you remove power to all system components. Any attached I/O enclosures power off completely. The Disk Drive enclosures are not placed in a low power state after the power off action or command.

To power off a system:

1. Stop applications. See [Stopping Application, Devices, and Processes](#) on page 15-14.
2. Use either of these methods to place most system components in a low-power state:
 - [System Power-Off Using OSM](#) on page 15-17
 - [System Power-Off Using SCF](#) on page 15-17
3. Shut off AC power to all peripheral devices.
4. Locate the circuit breaker that controls the power cords.
5. Switch the breakers off.

System Power-Off Using OSM

1. Log on to the OSM Service Connection.
2. Right-click the **System** object.
3. Select **actions**.
4. From the drop-down menu, select **System Power Off**.

System Power-Off Using SCF

To power off the system using SCF, log on to an available TACL command interpreter as the super ID (255,255) and issue the SCF power-off command:

```
> SCF CONTROL SUBSYS $ZZKRN, SHUTDOWN
```

Emergency Power-Off Procedure

If possible, HP recommends that the system be in a low power state before you remove power to the system. However, in emergency situations, you might need to quickly remove AC power from a system. Sites equipped with an emergency power-off (EPO) switch can use it to remove AC power from your entire system. For more information on the EPO system, see the *NonStop NS-Series Planning Guide*.

For sites that are not equipped with an EPO switch, switch off the circuit breakers to the power cords connected to the cabinets and peripherals. For more information, refer to [Section 14, Power Failures: Preparation and Recovery](#).

Troubleshooting and Recovery Operations

Refer to the appropriate subsection for recovery information:

- [Recovering From a System Load Failure](#) on page 15-20
- [Getting a Corrupt System Configuration File Analyzed](#) on page 15-21
- [Recovering From a Reload Failure](#) on page 15-21
- [Exiting the OSM Low-Level Link](#) on page 15-22
- [Opening Startup Event Stream and Startup TAACL Windows](#) on page 15-22

If any of these problems occur when you power on a system, see the appropriate subsection for recovery information:

- [Fans Are Not Turning](#) on page 15-18
- [System Does Not Appear to Be Powered On](#) on page 15-19
- [Green LED Is Not Lit After POSTs Finish](#) on page 15-19
- [Amber LED on a Component Remains Lit After the POST Finishes](#) on page 15-19
- [Components Fail When Testing the Power](#) on page 15-19

Fans Are Not Turning

If the fans do not start turning a few seconds after you power on the server, check that the AC power cords and component power cords are properly connected.

If the green LED lights are lit, but the fans are not turning, you must power off the system immediately. See in [Powering Off a System](#) on page 15-17. Contact your service provider.

System Does Not Appear to Be Powered On

If AC power is being supplied to the server and the fans are turning but the server still does not appear to be powered, the server might be running internal tests. Wait several minutes (at least 10 minutes for large configurations). If the server is still not powered on after this time and you cannot determine the cause of the problem:

- Check your site's circuit breakers.
- Plug in another device into the PDU that powers the LSU to check the power for that PDU.

Green LED Is Not Lit After POSTs Finish

It can take several minutes for the green LEDs on all system components to light:

1. Wait for the POSTs to finish. It might take as long as 10 minutes for all system components.
2. If the green LEDs still do not light:
 - a. Check that AC power cords and component power cords are properly connected.
 - b. If one green LED still does not light, a system component might have failed its POST.
3. If you cannot determine the cause of the problem, contact your service provider.

Amber LED on a Component Remains Lit After the POST Finishes

A fault might have been detected, or the component might not have been successfully initialized and configured. Contact your service provider.

Components Fail When Testing the Power

If a component fails when testing the power, the possible causes are listed in descending order of probability:

- The component is plugged in improperly.
Check the connection between each component power cord and PDU, and check the AC power receptacle to which the server power cord is connected. If necessary, unplug and plug each cord in again to ensure that it is seated properly.
- A power cord is defective.
- A component power supply is defective.
- A fuse is defective.
- The UPS, if installed, is not fully charged.

Recovering From a System Load Failure

If a system load is not successful or if the system halts:

1. Check for messages in the System Load dialog box of the OSM Low-Level Link.
2. Check the Processor Status dialog box of the OSM Low-Level Link for halt codes. Look up the halt codes in the *Processor Halt Codes Manual* for further information about the cause of failure and the appropriate recovery procedure.
3. Check the startup event message window for any event messages. Record the event messages, and refer to the appropriate documentation for recovery information. Refer to the *Operator Messages Manual* for further information about the cause, effect, and recovery procedure for an event.
4. Check that the disk you selected to load from is in the specified location.
5. Check that the disk is properly configured as a system disk.
6. Correct the problems, and try loading the system again. If you cannot correct the problem, contact your service provider.
7. If you continue to have problems, load the system from each disk path for both the primary and mirror \$SYSTEM drives.
8. If you cannot load the system using the current configuration file, load the system using a saved version of the system configuration file. See [Configuration File](#) on page 15-8.
9. If you still cannot load the system or if a CONF_{xyy} is not available, load the system from an alternate system disk if one is available.
10. If you cannot load the system from an alternate system disk, contact your service provider. You might be able to load the system from the CONFBASE configuration file and restore a previously backed-up configuration file. If you use this option, many additional steps are required to restore your system to normal working order.
 - a. Load the system as described in [Starting a System](#) on page 15-5. In the Configuration File box, select **Base (CONBASE)** as the configuration file.
 - b. Reload the remaining processors. See [Reloading Processors](#) on page 15-12.
 - c. From the Startup TAFL window, configure a tape drive.
 - d. Restore a previously backed-up configuration file.
 - e. Load the system as described in [Starting a System](#) on page 15-5 from the current configuration file (CONFIG). Check that the CIIN file is enabled.
11. After you load a system from a saved version (CONF_{xyy}) of the system configuration database file or the CONFBASE, verify that no pending changes to system attributes appear. From a TAFL prompt:

```
INFO SUBSYS $ZZKRN
```

Pending changes can appear (but are misleading) if the earlier configuration has different system name, number, or time attributes than the configuration you replaced. For example, if you load the \EAST system from the CONFBASE file (which specifies \NONAME as the system name), an INFO SUBSYS \$ZZKRN command displays \EAST as the current system and \NONAME as a pending change. Enter an ALTER SUBSYS command to change the system name to \EAST and cause the pending change to disappear. It is not displayed when you enter INFO.

Getting a Corrupt System Configuration File Analyzed

If the current system configuration file is corrupt, send it to your service provider for an analysis:

1. Return to a saved, stable configuration file. See [Configuration File](#) on page 15-8.
2. After the system is up and stable, copy to a backup tape the corrupt CONFSAVE file. For example:

```
> BACKUP $TAPE, $SYSTEM.ZSYSCONF.CONFSAVE, LISTALL
```

You must backup the CONFSAVE file before you perform the next system load. Another system load operation overwrites the CONFSAVE file you want analyzed.

3. Submit the tape to your service provider for analysis, along with a copy of any SCF command file or SCF log file of the commands that were part of the process that created the corrupt configuration.

Recovering From a Reload Failure

If a reload is not successful:

1. Check the Processor Status dialog box of the OSM Low-Level Link for halt codes. Look up the halt codes in the *Processor Halt Codes Manual* for further information about the cause of failure and the appropriate recovery procedure.
2. Check the System Load dialog box of the OSM Service Connection for messages.
3. Check for any event messages. Look up event messages in the EMS logs (\$0 and \$ZLOG) and refer to the OSM Event Viewer or the *Operator Messages Manual* for further information about the cause, effect, and recovery for any event message.
4. Perform a processor dump, if needed, as described in [Dumping a Processor to Disk](#) on page 9-15.
5. Try a soft reset of the processor.
6. Reload the processor or processors as described in [Section 9, Processors and Components: Monitoring and Recovery](#). If you cannot to correct the problem, contact your service provider.
7. If you still cannot reload the processor or processors, try a hard reset of the processor.

Exiting the OSM Low-Level Link

If all processors in the system have been halted and you are unable to log off, press Alt-F4 to exit the OSM Low-Level Link.

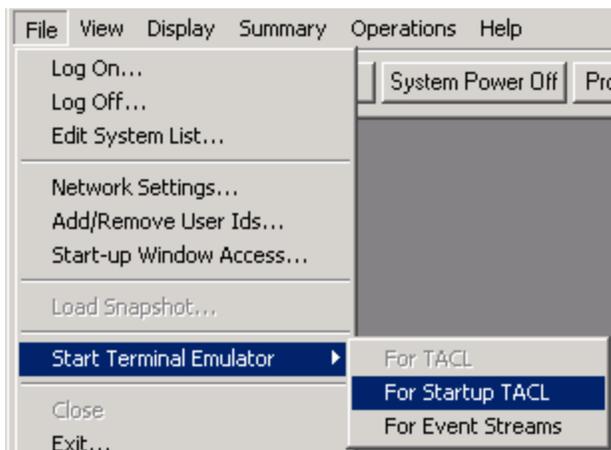
Opening Startup Event Stream and Startup TACL Windows

When you perform a normal system load, these windows open automatically. If the windows do not open or if you close them, you can reopen them using either OSM or Outside View.

To open startup event stream windows and startup TACL windows from OSM:

1. Log on to the OSM Low-Level Link.
2. From the File menu, select **Start Terminal Emulator > For Startup TACL**.

Figure 15-3. Opening a Startup TACL Window



VST813.vsd

3. Two OutsideView windows launch on top of the other. If you do not see the TACL prompt in one OutsideView window, check the other OutsideView window by clicking the buttons on the Windows toolbar.

Figure 15-4. OutsideView Buttons on the Windows Toolbar



4. From the File menu, select **Start Terminal Emulator > For Event Streams**.
5. Two OutsideView windows appear, but one launches on top of the other. If you do not see the TACL prompt in one OutsideView window, you can check the other OutsideView window (see [Figure 15-4](#)).

To open startup event stream windows and startup TACL windows using Outside View:

1. Select **Start > OutsideView > OutsideView**. The OutsideView dialog box appears.
2. Select **Session > New**. The New Session Properties dialog box appears.
3. On the Session tab, in the Session Caption box, type a session caption name such as Startup Events or Startup TACL.
4. Click **IO Properties**. The TCP/IP Properties dialog box appears.
5. Type the DNS name or IP address of the maintenance entity (ME) or integrated maintenance entity (IME), followed by a space, and the port number of the window type (303 for the startup event stream window, and 301 for the startup TACL window).

For example, for the startup event stream window:

```
me-system1- G100-M02 303
```

```
me-system1- G100-M03 303
```

```
192.231.36.2 303
```

Then, for the startup TACL window:

```
me-system1- G100-M02 301
```

```
me-system1- G100-M03 301
```

```
192.231.36.2 301
```

6. Click **OK**. The TCP/IP Properties dialog box closes, and you are returned to the New Session Properties dialog box.
7. Click **OK**. The startup event stream window or startup TACL window appears. A TACL prompt appears in the startup TACL window.

Consult the OSM Event Viewer for any important system startup messages logged to the EMS log that you might have missed while opening the terminal-emulation windows.

Related Reading

For more information, refer to the documentation listed in [Table 15-2](#).

Table 15-2. Related Reading for Starting and Stopping a System

For Information About	Refer to
Using SCF to customize your configuration	<i>SCF Reference Manual for H-Series RVUs</i> provides an overall reference for SCF, as well as information on customizing your configuration using command files.
Using TACL	<i>TACL Reference Manual</i>
Using the System Load dialog box to start or stop the system	OSM Low-Level Link online help
Creating an alternate system disk	<i>NonStop NS-Series Hardware Installation Guide</i>
Informing OSM of the location of an alternate system disk. See Saving (a disk-level action) or deleting (a system-level action) alternate system load volumes.	OSM Service Connection online help
Automating the startup of objects and devices using generic processes	<i>SCF Reference Manual for the Kernel Subsystem</i>
Starting storage devices such as disks and tape drives	<i>SCF Reference Manual for the Storage Subsystem</i>
Starting and stopping communications devices and communications lines	SCF or configuration manual specific to each type of communications device or line
Starting Ethernet addressable devices, including terminals and printers	The Configuration and Management manual for the communication subsystem for the terminal or printer
Starting WAN communications lines for devices and intersystem communications protocols	<i>WAN Subsystem Configuration and Management Manual</i> , as well as the SCF manuals that apply to the specific devices and communications protocols
Configuring the CIIN file	Section 16, Creating Startup and Shutdown Files <i>DSM/SCM User's Guide</i> for information on the CONFTEXT file CIIN entry
Shutting down DSM/SCM	<i>DSM/SCM User's Guide</i>
Starting up and shutting down TMF	<i>TMF Operations and Recovery Guide</i>
Starting and shutting down the Pathway environment	<i>Pathway/iTS System Management Manual</i> and <i>TS/MP System Management Manual</i>
Draining the spooler	<i>Spooler Utilities Reference Manual</i>

16

Creating Startup and Shutdown Files

This section describes command files that automatically start and shut down an Integrity NonStop NS-series server.

[Automating System Startup and Shutdown](#) on page 16-2

[Managed Configuration Services \(MCS\)](#) on page 16-2

[Startup](#) on page 16-2

[Shutdown](#) on page 16-3

[For More Information](#) on page 16-3

[Processes That Represent the System Console](#) on page 16-3

[\\$YMIOP.#CLCI](#) on page 16-3

[\\$YMIOP.#CNSL](#) on page 16-3

[\\$ZHOME](#) on page 16-4

[\\$ZHOME Alternative](#) on page 16-4

[Example Command Files](#) on page 16-4

[CIIN File](#) on page 16-5

[Establishing a CIIN File](#) on page 16-6

[Modifying a CIIN File](#) on page 16-6

[If a CIIN File Is Not Specified or Enabled in OSM](#) on page 16-7

[Example CIIN Files](#) on page 16-8

[Writing Efficient Startup and Shutdown Command Files](#) on page 16-9

[Command File Syntax](#) on page 16-9

[Avoid Manual Intervention](#) on page 16-10

[Use Parallel Processing](#) on page 16-10

[Investigate Product-Specific Techniques](#) on page 16-11

[How Process Persistence Affects Configuration and Startup](#) on page 16-11

[Tips for Startup Files](#) on page 16-11

[Startup File Examples](#) on page 16-12

[System Startup File](#) on page 16-12

[Spooler Warm-Start File](#) on page 16-14

[TMF Warm-Start File](#) on page 16-14

[TCP/IP Stack Configuration and Startup File](#) on page 16-14

[CP6100 Lines Startup File](#) on page 16-17

[ATP6100 Lines Startup File](#) on page 16-17

[X.25 Lines Startup File](#) on page 16-17

[Printer Line Startup File](#) on page 16-18

[Expand-Over-IP Line Startup File](#) on page 16-18

[Expand Direct-Connect Line Startup File](#) on page 16-18

[Tips for Shutdown Files](#) on page 16-19

[Shutdown File Examples](#) on page 16-19

[System Shutdown File](#) on page 16-20

[CP6100 Lines Shutdown File](#) on page 16-21

[ATP6100 Lines Shutdown File](#) on page 16-21

[X.25 Lines Shutdown File](#) on page 16-21

[Printer Line Shutdown File](#) on page 16-22

[Expand-Over-IP Line Shutdown File](#) on page 16-22

[Direct-Connect Line Shutdown File](#) on page 16-22

[Spooler Shutdown File](#) on page 16-23

[TMF Shutdown File](#) on page 16-23

Automating System Startup and Shutdown

Managed Configuration Services (MCS)

Integrity NonStop NS-Series servers are being configured with Managed Configuration Services (MCS) when they are installed. Functions performed by MCS include configuring the Spooler, Pathway, and Expand lines, and creating startup and shutdown files. This section is about creating startup and shutdown files without the use of MCS. For more information about MCS, see the *Integrity NonStop NS-Series Hardware Installation Manual*.

Startup

You can use startup command files to automate the starting of devices and processes on the system, which minimizes the possibility of operator errors caused by forgotten or mistyped commands.

The system is shipped with a basic startup file named CIIN, located on the \$SYSTEM.SYS00 subvolume. The CIIN file must be specified in a particular way. See [CIIN File](#) on page 16-5 for more information.

After the commands in the CIIN file are executed, other startup files can be invoked either automatically, from another startup file, or manually in commands entered by the operator. The startup file sequence usually starts the spooler and other system software first and then starts applications.

Shutdown

Automating system shutdown with a collection of shutdown files helps the operator bring the system to an orderly halt. The shutdown file sequence reverses the order of commands in the startup file sequence: applications are shut down first, followed by the spooler and other system software.

For More Information

For information about ...	See
Configuring the spoolers	<i>Spooler Utilities Reference Manual</i>
Configuring and managing Pathway applications	<i>TS/MP System Management Manual</i>
Configuring and managing TMF	<i>TMF Planning and Configuration Guide</i> <i>TMF Operations and Recovery Guide</i>
Configuring and managing TRANSFER applications	<i>TRANSFER Installation and Management Guide</i>

Processes That Represent the System Console

On Integrity NonStop NS-series servers, the system console is a pair of windows on a LAN-connected system console. It is represented by the processes \$YMIOP.#CLCI and \$YMIOP.#CNSL, and the home terminal is represented by the \$ZHOME process pair.

\$YMIOP.#CLCI

\$YMIOP.#CLCI is the primary interactive terminal for the operator interface to the system.

This process:

- Runs on the system console
- Is preconfigured on your system during system generation

TACL processes are started on \$YMIOP.#CLCI by commands in the CIIN file. If a read operation is pending, such as a TACL prompt on \$YMIOP.#CLCI, write operations are blocked, causing the process attempting the write operation to wait indefinitely.

\$YMIOP.#CNSL

\$YMIOP.#CNSL is a write-only device for logging.

This process:

- Runs on the system console

- Is preconfigured on your system during system generation

\$ZHOME

The \$ZHOME process is a process pair that provides a reliable home terminal to which processes can perform write operations. The \$ZHOME process can be used by processes that must write to the system console but do not require a response.

\$ZHOME is preconfigured on your system by the CONFBASE file. \$ZHOME is a generic process that is part of the SCF Kernel subsystem. Note the following about the configuration of \$ZHOME:

- The \$ZHOME process is configured with \$YMIOP.#CLCI as its HOMETERM, INFILE, and OUTFILE.
- Because \$ZHOME acts as a reliable home terminal designed to interact with the system console, \$YMIOP.#CLCI, HP recommends that you do not change its configuration. Most important:
 - Do not specify \$ZHOME for the INFILE, OUTFILE, or HOMETERM for the \$ZHOME process.
 - Never specify \$ZHOME for the INFILE for a process. The \$ZHOME process returns the FEINVALOP error (file-system error 2) in response to any read operation.
- Generic processes started by the \$ZPM persistence manager inherit \$YMIOP.#CLCI as the HOMETERM, INFILE, and OUTFILE unless these attributes are changed in the configuration record for the generic process. HP recommends that you configure most NonStop Kernel and system-level generic processes to use \$ZHOME for the HOMETERM and OUTFILE.

\$ZHOME Alternative

Instead of \$ZHOME, you might want to use the optional NonStop Virtual Hometerm Subsystem (VHS) product if both of the following conditions are true:

- The processes you are configuring cannot handle error responses returned if \$YMIOP.#CNSL or \$YMIOP.#CLCI is not available.
- The process must perform read operations to the device.

Example Command Files

This section describes and shows examples of command files that can be used to start up and shut down the server:

- Examples and sample programs are for illustration only and might not be suited for your particular purpose.

HP does not warrant, guarantee, or make any representations regarding the use or the results of the use of any examples or sample programs in any documentation. You must verify the applicability of any example or sample program before placing the software into production use.

- These examples are for a system whose configuration has been changed from the factory-installed configuration. Your system's initial configuration will differ from these examples.

The startup files in this section assume that the objects they start have already been added to the system configuration database.

- The IP addresses used in this section are examples only. If you use the example files described in this section on your system, you must change the IP addresses in these examples to IP addresses that are appropriate for your LAN environment.
- The configuration track-ID for the SWAN concentrator used in the example files, X001XX, is also an example.

If you use the example files described in this section on your system, you must change the configuration track-ID used in these examples to the actual configuration track-ID assigned to your SWAN concentrator.

CIIN File

The CIIN file is a TACL command (OBEY) file that contains a limited set of commands that usually:

- Start a TACL process pair on the system console for the system console TACL window (\$YMIOP.#CLCI)

When the startup TACL executes the commands in the CIIN file and terminates, the \$YMIOP.#CLCI process pair lets you log on to the system and complete the system startup.

Note. Before these TACL processes start, open the appropriate terminal emulator windows with the OSM Low-Level Link. You must open these windows before performing a system load.

OSM software lets you define primary and backup IP addresses for TACL windows. For more information about configuring OSM software, see the *OSM User's Guide*.

- Load all processors that are not currently running

Alternatively, the CIIN file can reload a minimal set of processors, such as processor 1, to bring up a minimal system. You can then test for successful startup of a minimal system environment before you bring up the remainder of the system.

Normally, the initial TACL process invokes the CIIN file automatically after the first processor is loaded if all these conditions are true:

- The CONFTEXT configuration file located in the \$SYSTEM.SYS_{nn} subvolume has an INITIAL_COMMAND_FILE entry for the CIIN file.
- The CIIN file is available in the specified location.
- The CIIN option is not disabled in the System Load dialog box.

Note. By default, the CIIN file contains commands needed to start the (permanent) TACL process pair and to reload all the processors in the system. Do not place commands to prime the processors in the CIIN file.

Establishing a CIIN File

The CIIN file is configured at the factory as \$DSMSCM.SYS.CIIN. You do not need to establish this file. DSM/SCM automatically copies the CIIN file from the initial location into each SYS_{nn} you create.

Note. The CIIN file must be owned by a member of the super-group (255,_n). HP recommends that you specify “N” for the read access portion of the file security attribute (RWEF) to allow the file to be read by any user on the network. For example, you might secure this file “NCCC.”

The name of the CIIN file is specified in the INITIAL_COMMAND_FILE entry of the CONFTEXT configuration file.

A system generation program (run from the DSM/SCM application) copies the file specified in the CONFTEXT file onto the SYS_{nn} subvolume on the disk and renames the file CIIN. If no file is specified in CONFTEXT, the operating system does not look for the startup file SYS_{nn}.CIIN at system startup, even if you enable that file. You cannot simply copy a startup file to the SYS_{nn} subvolume and name it CIIN.

Modifying a CIIN File

After the CIIN file is established on \$SYSTEM.SYS_{nn} (as part of running DSM/SCM), you can modify the contents of SYS_{nn}.CIIN with a text editor such as TEDIT. You need not run DSM/SCM again to make these changes effective.

If a CIIN File Is Not Specified or Enabled in OSM

The results of the startup TACL process varies depending on whether a CIIN file is specified in the CONFTEXT file and whether the CIIN option is enabled.

CONFTEXT CIIN Entry and CIIN File	CIIN Option	Results
1. CONFTEXT has CIIN entry, and file is available in specified location.	Enabled	CIIN is executed by initial (startup) TACL process. Upon completion, this TACL process terminates. You must log on to a different TACL process (the TACL process on \$YMIOP.#CLCI started by the CIIN file) to complete the system startup process.
2. CONFTEXT has CIIN entry, and file is available in specified location, but file is empty or aborts because of syntax errors before another TACL process is started.	Enabled	CIIN is executed by initial (startup) TACL process. Upon completion, this TACL process terminates, leaving no TACL process available. You must reload the system with the CIIN option disabled in the System Startup dialog box invoked from the OSM Low-Level Link; then log on and correct the CIIN file. Then, either enable the CIIN option, using the System Startup dialog box, and reload; or complete the system startup process manually.
3. CONFTEXT has CIIN entry, but file is not available in specified location.	Enabled	Initial TACL process is started and left in logged-off state. You must log on to complete the system startup process.
4. CONFTEXT has CIIN entry.	Disabled	Initial TACL process is started and left in logged-off state. You must log on to complete the system startup process.
5. CONFTEXT has no CIIN entry.	Enabled or disabled	Initial TACL process is started and left logged on to the super ID (255,255). You must initiate the remainder of the system startup process manually and then log off.

△ Caution. Situation 5 presents a security issue: the initial TACL process is left logged on to the super ID (255,255). You must either immediately continue with the system startup process (as described in the Results column), log on to another user ID, or log off.

Example CIIN Files

This example CIIN file does not include a persistent CLCI TACL process.

```
Comment -- This is the initial command input (CIIN) file for the system.
Comment -- If CIIN is enabled in OSM and configured in your CONFTEXT
Comment -- file, the initial TACL process will read this file and
Comment -- then terminate.
```

```
Comment -- This file is used to reload the remaining processors and
Comment -- start a TACL process pair for the system console.
```

```
Comment -- Reload the remaining processors.
```

```
RELOAD /TERM $ZHOME, OUT $ZHOME/ *
```

```
Comment -- Start a TACL process pair for the system console TACL window.
Comment -- Use the OSM Low-Level Link to start a TTE session
Comment -- for the startup TACL before issuing this command (see the
Comment -- Start Terminal Emulator command under the File menu).
Comment -- These should be the last commands in this file, because
Comment -- the TACL process displays a prompt and attempts to read
Comment -- from $YMIOP.#CLCI, blocking other processes from writing to
Comment -- this device.
```

```
TACL/TERM $YMIOP.#CLCI, IN $YMIOP.#CLCI, OUT $YMIOP.#CLCI, NAME $SC0, &
PRI 199, CPU 0/1
```

```
TACL/TERM $YMIOP.#CLCI, IN $YMIOP.#CLCI, OUT $YMIOP.#CLCI, NAME $SC0, &
PRI 199, CPU 1/0
```

```
Comment -- Upon completion of this file, the initial TACL process
Comment -- terminates. You need to log on to a new TACL session
Comment -- to complete the remainder of the system startup process.
```

This example CIIN file shows what you would use if you had created a persistent CLCI TACL process by configuring it as a generic process.

-
- △ Caution.** If you use the following example as *is*, you must first create the \$ZZKRN.#CLCI-TACL process. If you do not create the \$ZZKRN.#CLCI-TACL process first, you might have no access to the system after it loads. To recover, load the system again from another SYS_{nn} or with CIIN disabled.
-

```
Comment -- This is the initial command input (CIIN) file for the system.
Comment -- If CIIN is enabled in OSM and configured in your CONFTEXT
Comment -- file, the initial TACL process will read this file and
Comment -- then terminate.
```

```
Comment -- This file is used to reload the remaining processors and
Comment -- start a TACL process pair for the system console.
```

```
Comment -- Reload the remaining processors.
```

```
RELOAD /TERM $ZHOME, OUT $ZHOME/ *
```

```
Comment -- Use SCF to start a persistent TACL process pair for the
Comment -- system console TACL window.
Comment -- Use the OSM Low-Level Link to start a TTE session
Comment -- for the startup TACL before issuing this command (see the
Comment -- Start Terminal Emulator command under the File menu). This SCF
Comment -- command must be the last command in this file, because the TACL
Comment -- process creates displays a prompt and attempts to read from
Comment -- $YMIOP.#CLCI, blocking other processes from writing to this
Comment -- device.
```

```
SCF /NOWAIT, OUT/ START PROCESS $ZZKRN.#CLCI-TACL
```

Writing Efficient Startup and Shutdown Command Files

TACL and by many subsystems support command files. Command files for startup or shutdown contain a series of commands that automatically execute when the file is executed. To automate and reduce the time required to start and stop your applications, devices, and processes:

- Include commands in one or more command files that you invoke from either a TACL prompt or another file.
- Write efficient startup and shutdown command files.
 - Use command file syntax that executes quickly.
 - Avoid manual intervention to ensure that command files execute quickly
 - Use parallel processing to distribute startup and shutdown processes across multiple processors.
 - Investigate and use product-specific techniques for fast startup and shutdown.

Command File Syntax

The syntax in command files affects the time it takes for them to execute. To ensure that your command files execute quickly:

- Avoid using wild-card characters in command files

A wild card is a character—typically an asterisk (*) or a question mark (?)—used to match any character or series of characters. When you use wild-card characters in your command files, execution time is increased because the system must look up names in a table. By using explicit names instead of wild-card characters, you shorten execution time and allow for commands to execute in parallel.

This PATHCOM START command uses a wild-card character to start all of the TERM objects defined in the PATHMON configuration file:

```
= START TERM *
```

This PATHCOM START command uses explicit names to start all of the TERM objects defined in the PATHMON configuration file:

```
= START TERM (TERM1, TERM2, TERM3, TERM4, TERM5, TERM6)
```

Note. When using explicit names, you must revise your command files whenever a configuration change occurs. Therefore, you should balance the time it takes to update configuration files against the savings in startup or shutdown time.

- Use single-line commands instead of multiple-line commands. Multiple-line commands in a command file increase execution time.

Avoid Manual Intervention

Write startup and shutdown files so that they execute correctly without requiring manual intervention. Any time an operator must intervene, startup and shutdown time increase and the possibility of human error increases.

Use Parallel Processing

Parallel processing decreases the time required to start up or shut down your system or application because startup and shutdown processes are distributed throughout the processors in your system. For example, this SCF command file uses parallel processing in four processors to start several communications lines. The files START0, START1, START2, and START3 contain the actual commands that start the communications lines.

This command file uses a special technique intended to ensure that each process gets started even if a given processor is out of service. The technique is to start each process in two processors. If the first processor is down, the command file continues to the next processor. If the first processor is up, and the process is started, the command file still continues to the next processor but fails because the process name (SS_n) is in use by the process that was successfully started. As a result, a specified process is started in whichever processor is running. Of course, if neither processor is up, the attempt to start the process fails.

```
SCF /IN START0, NOWAIT, CPU 0, NAME $S0/
```

```
SCF /IN START0, NOWAIT, CPU 2, NAME $S0/
```

```
SCF /IN START1, NOWAIT, CPU 1, NAME $S1/
```

```
SCF /IN START1, NOWAIT, CPU 3, NAME $S1/
```

```
SCF /IN START2, NOWAIT, CPU 2, NAME $S2/
```

```
SCF /IN START2, NOWAIT, CPU 0, NAME $S2/
```

```
SCF /IN START3, NOWAIT, CPU 3, NAME $S3/
```

```
SCF /IN START3, NOWAIT, CPU 1, NAME $S3/
```

When using the technique shown in this command file, make sure to spread the process workload across all available processors. If there are too many processes to start in processors 0 and 1, queuing and memory-contention problems can result.

Investigate Product-Specific Techniques

Some products provide commands that reduce the time required to start up or shut down their services. Familiarize yourself with the products and applications that run on your system to identify time-saving techniques for speeding startup and shutdown operations. Refer to the relevant documentation for each product.

For example, the HP NonStop TS/MP product provides the COOL START option and the SHUTDOWN2 command to shorten startup and shutdown times, respectively. Using the COOL START option, rather than COLD START, to restart an existing transaction-processing system is much faster. The SHUTDOWN2 command is faster and more reliable than the SHUTDOWN command. Both of these techniques are described in the *TS/MP System Management Manual*.

How Process Persistence Affects Configuration and Startup

When the system is started, all processes that are configured to be persistent are started automatically by the persistence manager (\$ZPM) or by the subsystem manager, which is started by \$ZPM.

For example, when the system is started, the WAN subsystem manager automatically starts all WAN I/O processes (IOPs) that were started before the system was shut down. However, communications lines and paths must be started manually by the operator.

To make important system processes start automatically at system load and be persistent (that is, restart automatically if stopped abnormally), you should create them as generic processes in the system configuration database. See the *Integrity NonStop NS-Series Hardware Installation Manual*.

For more information about persistence and the \$ZPM persistence manager, see the *SCF Reference Manual for the Kernel Subsystem*.

Tips for Startup Files

HP recommends that you specify “N” for the read access portion of the file security attribute (RWEF) for your startup files to allow the files to be read by any user on the network. For example, you might secure these files “NCCC.”

The sequence in which you invoke startup files can be important. Some processes require other processes to be running before they can be started. Be sure to indicate the order in which your startup files are to be run.

Because the TCP/IP configurations are not stored in the configuration database, they are not preserved after system loads. Therefore, TCP/IP stacks must be configured as well as started each time the system is started. This is only true for conventional TCP/IP.

Startup File Examples

You can implement the system startup sequence with a collection of startup files, each with a specific purpose. HP recommends that you invoke the startup files in this order:

1. Startup file for the system, to be invoked after the CIIN file is invoked
2. Startup files for the system software
3. Startup files for the subsystems
4. Startup files for the communications lines
5. Startup files for the applications

See [Section 15, Starting and Stopping the System](#) for detailed instructions on the startup procedure. For information about automating disk processes upon startup, see the *Integrity NonStop NS-Series Planning Guide*.

Note. Examples and sample programs are for illustration only and might not be suited for your particular purpose. HP does not warrant, guarantee, or make any representations regarding the use or the results of the use of any examples or sample programs in any documentation. You must verify the applicability of any example or sample program before placing the software into production use. For more information, see [Example Command Files](#) on page 16-4.

System Startup File

The following example shows a partial command file that starts up the system software and invokes other startup files.

After the commands in the CIIN file have been executed and the initial system startup sequence is complete, the local operator invokes this file by entering the following TACL command:

```
> OBEY $SYSTEM.STARTUP.STRTSYS

Comment -- This is $SYSTEM.STARTUP.STRTSYS

comment -- Start the server for labeled tape processing.

ZSERVER / NAME $ZSVR, NOWAIT, PRI 145, CPU 0 / 1
ZSERVER / NAME $ZSVR, NOWAIT, PRI 145, CPU 1 / 0
MEDIACOM ALTER TAPEDRIVE *, NLCHECK OFF

Comment -- If you have used SCF to start a persistent Subsystem
Comment -- Control Process (SCP) process pair, you do not need an
Comment -- explicit SCP command to start $ZNET, unless you load the
Comment -- system from a different CONFIG file.
Comment -- All SCF commands are routed through the SCP process.  $ZNET
Comment -- routes each request to the appropriate communication
Comment -- management process (such as Expand or SNAX).

Comment -- If you have not configured SCP as a persistent generic
Comment -- process, remove the commenting from the following SCP
Comment -- command and start SCP as a nonpersistent process pair.
Comment -- SCP / NAME $ZNET, NOWAIT, PRI 199, TERM $ZHOME, OUT $ZHOME, &
Comment -- CPU 0/1; AUTOSTOP -1
```

```
Comment -- If you have used SCF to start a persistent $ZEXP Expand
Comment -- manager process pair, you do not need an explicit SCP
Comment -- command to start $ZEXP, unless you load the system from a
Comment -- different CONFIG file.

Comment -- If you have not configured $ZEXP as a persistent generic
Comment -- process, remove the commenting from the following SCP
Comment -- command and start $ZEXP as a nonpersistent process pair.
Comment -- OZEXP / NAME $ZEXP, NOWAIT, PRI 180, OUT $ZHOME, CPU 0/1

comment -- Warm start the spooler subsystem using the SPOOLCOM command
comment -- file SPLWARM

OBEY $SYSTEM.STARTUP.SPLWARM

comment -- Start the Transaction Management Facility (TMF) subsystem
comment -- using the TMFCOM command file TMFSTART

TMFCOM / IN $SYSTEM.STARTUP.TMFSTART, OUT $ZHOME /

comment -- Configure and start the TCP/IP stacks on the LAN adapter ports
comment -- used by the SWAN

OBEY $SYSTEM.STARTUP.IPSTK*

comment -- Start the CP6100 lines on the SWAN

SCF / IN $SYSTEM.STARTUP.STRTCP6, OUT $ZHOME /

comment -- Start the ATP6100 lines on the SWAN

SCF / IN $SYSTEM.STARTUP.STRTATP, OUT $ZHOME /

comment -- Start the X.25 lines on the SWAN

SCF / IN $SYSTEM.STARTUP.STRTX25, OUT $ZHOME /

comment -- Start the printers on the SWAN

SCF / IN $SYSTEM.STARTUP.STRTLTP, OUT $ZHOME /

comment -- Start the Expand-over-IP line to \Case2

SCF / IN $SYSTEM.STARTUP.IP2CASE2, OUT $ZHOME /

comment -- Start the direct-connect line

SCF / IN $SYSTEM.STARTUP.STRTLH, OUT $ZHOME /
```

Spooler Warm-Start File

This example command file warm starts the spooler.

After the spooler has been brought up, the printer devices should be in the WAITING state. This file can be invoked automatically from the STRTSYS file, or you can invoke it by using the following TACL command:

```
> OBEY $SYSTEM.STARTUP.SPLWARM

comment -- This is $SYSTEM.STARTUP.SPLWARM
comment -- This file warm starts the spooler, leaving all jobs intact.

SPOOL / IN $SYSTEM.SPL.SPL, OUT $ZHOME, NAME $SPLS, NOWAIT, PRI 149, &
CPU 1/0
SPOOLCOM; SPOOLER, START

comment -- check to see that the spooler started successfully

SPOOLCOM; SPOOLER, STATUS
```

TMF Warm-Start File

This example command file warm starts the TMF subsystem.

This file can be invoked automatically from the STRTSYS file, or you can invoke it by using the following TACL command:

```
> TMFCOM / IN $SYSTEM.STARTUP.TMFSTART, OUT $ZHOME /

-- This is $SYSTEM.STARTUP.TMFSTART
-- This file warm starts the Transaction Management Facility (TMF) subsystem
-- and checks to see if TMF started successfully.

START TMF;ENABLE DATAVOLS *;STATUS TMF;EXIT
```

TCP/IP Stack Configuration and Startup File

Configuration data for NonStop TCP/IP (conventional TCP/IP) processes is not added to the configuration database. Therefore, TCP/IP stacks must be both configured and started for each LAN port that connects to a SWAN concentrator each time you start the system, unless you are using NonStop TCP/IPv6 over SWAN. (If so, see the manuals that support those TCP/IP subsystems.)

You can create TACL command files to configure TCP/IP stacks on the other ports by assigning the appropriate values the following variables in this example file:

GW^ADDR	LINE^NAME	TCP^CPU2
HOST^NAME	LST^NAME	TCP^NAME
IP^ADDR	TCP^CPU1	TEL^NAME

The NonStop TCP/IPv6 subsystems participate in the system configuration database (however, not with the initial configuration database that is shipped with a new system).

- *TCP/IP Configuration and Management Manual*
- *TCP/IPv6 Configuration and Management Manual*

This example shows a TACL command file that configures the TCP/IP stack on \$ZZLAN.L018.

```
?TACL MACRO

== This file is $SYSTEM.STARTUP.IPSTK1
== Adds TCPIP and related processes to $ZZLAN.L018

#FRAME
#PUSH CON^NAME, LINE^NAME, TCP^NAME, LST^NAME, TEL^NAME
#PUSH HOST^NAME, IP^ADDR, GW^ADDR, TCP^CPU1, TCP^CPU2

#SET IP^ADDR    192.231.36.099
#SET GW^ADDR    192.231.36.17

#SET CON^NAME   $ZHOME
#SET LINE^NAME  L018
#SET TCP^NAME   $ZB018
#SET LST^NAME   $ZP018
#SET TEL^NAME   $ZN018
#SET HOST^NAME  "Case1_L018.DevInc.com"
#SET TCP^CPU1   0
#SET TCP^CPU2   1

[ #IF NOT [ #PROCESSEXISTS $ZNET ]
  | THEN |
  #OUTPUT
  #OUTPUT Starting SCP...
  SCP /NAME $ZNET, NOWAIT, CPU 0, PRI 165, TERM [CON^NAME]/ 1; AUTOSTOP -1
]

[ #IF [ #PROCESSEXISTS [LST^NAME] ]
  | THEN |
  STOP [LST^NAME]
]

#OUTPUT

#OUTPUT Stopping existing TCP/IP processes...
[ #IF [ #PROCESSEXISTS [TEL^NAME] ]
  | THEN |
  STOP [TEL^NAME]
]

[ #IF [ #PROCESSEXISTS [LST^NAME] ]
  | THEN |
  STOP [LST^NAME]
]

[ #IF [ #PROCESSEXISTS [TCP^NAME] ]
  | THEN |
  #PUSH #INLINEPREFIX
  SET VARIABLE #INLINEPREFIX +
  SCF /INLINE, OUT [#MYTERM], NAME/
  + ALLOW ALL ERRORS
  + ABORT PROCESS [TCP^NAME]
```

```

    + EXIT
    #POP #INLINEPREFIX
]
#OUTPUT
#OUTPUT Starting TCP/IP: [TCP^NAME]
TCPIP /NAME [TCP^NAME], TERM [CON^NAME], NOWAIT, CPU [TCP^CPU1] / [TCP^CPU2]
DELETE DEFINE =TCPIP^PROCESS^NAME
ADD     DEFINE =TCPIP^PROCESS^NAME, FILE [TCP^NAME]
PARAM TCPIP^PROCESS^NAME [TCP^NAME]
PARAM ZTNT^TRANSPORT^PROCESS^NAME [TCP^NAME]

#OUTPUT
#OUTPUT Configuring TCP/IP...
PUSH #INLINEPREFIX
SET VARIABLE #INLINEPREFIX +
SCF /INLINE, OUT [#MYTERM], NAME/
+ ALLOW ALL ERRORS
+ ASSUME PROCESS [TCP^NAME]
+ ALTER , HOSTNAME [HOST^NAME]
+ ADD SUBNET #SN1, TYPE ETHERNET, IPADDRESS [IP^ADDR], DEVICENAME [LINE^NAME]
+ ALTER SUBNET #SN1, SUBNETMASK %%hFFFFFFF00
+ ALTER SUBNET #LOOP0, IPADDRESS 127.1
+ START SUBNET *
+ ADD ROUTE #GW, DESTINATION 0, GATEWAY [GW^ADDR], DESTTYPE BROADCAST
+ START ROUTE *
+ EXIT
POP #INLINEPREFIX

#OUTPUT
#OUTPUT Starting Listener: [LST^NAME]
LISTNER /NAME [LST^NAME], CPU [TCP^CPU1], PRI 160, NOWAIT, TERM [CON^NAME],
HIGHPIN OFF/ $SYSTEM.ZTCPIP.PORTCONF

#OUTPUT
#OUTPUT Starting Telnet: [TEL^NAME]
TELSERV /NAME [TEL^NAME], CPU [TCP^CPU1], PRI 170, NOWAIT, TERM [CON^NAME]/ -
backupcpu [TCP^CPU2]

#OUTPUT
#OUTPUT Starting Telnet: [TEL^NAME]
TELSERV /NAME [TEL^NAME], CPU [TCP^CPU1], PRI 170, NOWAIT, TERM [CON^NAME]/ -
backupcpu [TCP^CPU2]

DELETE DEFINE =TCPIP^PROCESS^NAME
CLEAR  PARAM TCPIP^PROCESS^NAME
CLEAR  PARAM ZTNT^TRANSPORT^PROCESS^NAME
#UNFRAME

```

CP6100 Lines Startup File

This example shows an SCF command file that starts the CP6100 lines associated with the SWAN concentrator \$ZZWAN.#S01 (configuration track-ID X001XX).

This file can be invoked automatically from the STRTSYS file, or you can invoke it by using the following TACL command:

```
> SCF / IN $SYSTEM.STARTUP.STRTCP6, OUT $ZHOME /
== This is $SYSTEM.STARTUP.STRTCP6
== Starts CP6100 lines associated with the SWAN concentrator
== $ZZWAN.#S01
ALLOW 20 ERRORS
START LINE $CP6*
```

ATP6100 Lines Startup File

This example shows an SCF command file that starts the ATP6100 lines associated with the SWAN concentrator \$ZZWAN.#S01 (configuration track-ID X001XX).

This file can be invoked automatically from the STRTSYS file, or you can invoke it by using the following TACL command:

```
> SCF / IN $SYSTEM.STARTUP.STRTATP, OUT $ZHOME /
== This is $SYSTEM.STARTUP.STRTATP
== Starts ATP6100 lines associated with the SWAN concentrator
== $ZZWAN.#S01
ALLOW 20 ERRORS
START LINE $ATP*
```

X.25 Lines Startup File

This example shows an SCF command file that starts the X.25 lines associated with the SWAN concentrator \$ZZWAN.#S01 (configuration track-ID X001XX).

This file can be invoked automatically from the STRTSYS file, or you can invoke it by using the following TACL command:

```
> SCF / IN $SYSTEM.STARTUP.STRTX25, OUT $ZHOME /
== This is $SYSTEM.STARTUP.STRTX25
== Starts the X.25 lines associated with the SWAN concentrator
== $ZZWAN.#S01
ALLOW 20 ERRORS
START LINE $X25*
```

Printer Line Startup File

This example shows an SCF command file that starts a printer line associated with the SWAN concentrator \$ZZWAN.#S01 (configuration track-ID X001XX).

This file can be invoked automatically from the STRTSYS file, or you can invoke it by using the following TACL command:

```
> SCF / IN $SYSTEM.STARTUP.STRTLP, OUT $ZHOME /
== This is $SYSTEM.STARTUP.STRTLP
== Starts the printer associated with the SWAN concentrator
== $ZZWAN.#S01
ALLOW 20 ERRORS
START LINE $LP5516
```

Expand-Over-IP Line Startup File

This example shows an SCF command file that starts an Expand-over-IP communications line from \$ZZLAN.LAN08, at IP address 192.231.36.094, to \Case2, a NonStop K-series server at IP address 192.231.36.089.

This file can be invoked automatically from the STRTSYS file, or you can invoke it by using the following TACL command:

```
> SCF / IN $SYSTEM.STARTUP.IP2CASE2, OUT $ZHOME /
```

Note that the IP addresses used in this file are examples only. If you use this example file on your system, you must change these IP addresses to IP addresses that are appropriate for your LAN environment.

```
== This is $SYSTEM.STARTUP.IP2CASE2
ALLOW 100 ERRORS

START LINE $Case2IP
```

Expand Direct-Connect Line Startup File

This example shows an SCF command file that starts an Expand direct-connect line on a SWAN concentrator.

This file can be invoked automatically from the STRTSYS file, or you can invoke it by using the following TACL command:

```
> SCF / IN $SYSTEM.STARTUP.STRTLH, OUT $ZHOME /
== This is $SYSTEM.STARTUP.STRTLH
START LINE $Case2elh
```

Tips for Shutdown Files

HP recommends that you specify “N” for the read access portion of the file security attribute (RWEF) for your shutdown files to allow the files to be read by any user on the network. For example, you might secure these files “NCCC.”

The sequence in which you invoke shutdown files can be important. Some processes require other processes to be stopped before they can be stopped. Be sure to indicate the order in which shutdown files are to be run.

Shutdown File Examples

You can implement the system shutdown sequence with a collection of shutdown files, each with a specific purpose. HP recommends that you invoke the shutdown files in this order:

1. Shutdown files for the applications
2. Shutdown files for the communications lines
3. Shutdown files for the subsystems
4. Shutdown files for the system software
5. Shutdown file for the system

Note. Examples and sample programs are for illustration only and might not be suited for your particular purpose. HP does not warrant, guarantee, or make any representations regarding the use or the results of the use of any examples or sample programs in any documentation. You must verify the applicability of any example or sample program before placing the software into production use. For other information about these examples, see [Example Command Files](#) on page 16-4.

System Shutdown File

This example shows a TACL command file that shuts down the system software and invokes other shutdown files.

The local operator invokes this file by entering the following TACL command:

```
> OBEY $SYSTEM.SHUTDOWN.STOPSYS
```

Note. Shutting down the system in an orderly fashion does not require that you shut down every process. Some processes that have startup files might not need shutdown files.

```
comment -- This is $SYSTEM.SHUTDOWN.STOPSYS
comment -- Use this file to shut the system down in an orderly fashion.
comment -- Shut down the CP6100 lines associated with the SWAN concentrator
SCF/ IN $SYSTEM.SHUTDOWN.SDNCP6, OUT $ZHOME /

comment -- Shut down the ATP6100 lines associated with the SWAN concentrator
SCF/ IN $SYSTEM.SHUTDOWN.SDNATP, OUT $ZHOME /

comment -- Shut down the X.25 lines associated with the SWAN concentrator
SCF/ IN $SYSTEM.SHUTDOWN.SDNX25, OUT $ZHOME /

comment -- Shut down the printer lines associated with the SWAN concentrator
SCF/ IN $SYSTEM.SHUTDOWN.SDNLP, OUT $ZHOME /

comment -- Shut down the Expand-over-IP line to \Case2
SCF/ IN $SYSTEM.SHUTDOWN.IP2CASE2, OUT $ZHOME /

comment -- Shut down the Expand manager process, $ZEXP
SCF/ IN $SYSTEM.SHUTDOWN.SDNEXP, OUT $ZHOME /

comment -- Shut down the direct-connect line
SCF/ IN $SYSTEM.SHUTDOWN.STRTLH, OUT $ZHOME /

comment -- Drain the spooler subsystem using the SPOOLCOM command file
comment -- SPLDRAIN.

OBEY $SYSTEM.SHUTDOWN.SPLDRAIN

comment -- Stop the Transaction Management Facility (TMF) subsystem using the
comment -- TMFCOM command file TMFSTOP.

TMFCOM / IN $SYSTEM.SHUTDOWN.TMFSTOP, OUT $ZHOME /
```

CP6100 Lines Shutdown File

This example shows an SCF command file that stops the ATP6100 lines associated with the SWAN concentrator \$ZZWAN.#S01 (configuration track-ID X001XX).

This file can be invoked automatically from the STOPSYS file, or you can invoke it by using the following TACL command:

```
> SCF/ IN $SYSTEM.SHUTDOWN.SDNCP6, OUT $ZHOME /
== This is $SYSTEM.SHUTDOWN.SDNCP6
== This shuts down the CP6100 lines associated with the SWAN concentrator
== $ZZWAN.#S01
ALLOW 20 ERRORS
ABORT LINE $cp6*
```

ATP6100 Lines Shutdown File

This example shows an SCF command file that stops the ATP6100 lines associated with the SWAN concentrator \$ZZWAN.#S01 (configuration track-ID X001XX).

This file can be invoked automatically from the STOPSYS file, or you can invoke it by using the following TACL command:

```
> SCF/ IN $SYSTEM.SHUTDOWN.SDNATP, OUT $ZHOME /
== This is $SYSTEM.SHUTDOWN.SDNATP
== This shuts down the ATP6100 lines associated with the SWAN concentrator
== $ZZWAN.#S01
ALLOW 20 ERRORS
ABORT LINE $ATP*
```

X.25 Lines Shutdown File

This example shows an SCF command file that stops the X.25 lines associated with the SWAN concentrator \$ZZWAN.#S01 (configuration track-ID X001XX).

This file can be invoked automatically from the STOPSYS file, or you can invoke it by using the following TACL command:

```
> SCF/ IN $SYSTEM.SHUTDOWN.SDNX25, OUT $ZHOME /
== This is $SYSTEM.SHUTDOWN.SDNX25
== This shuts down the X.25 lines associated with the SWAN concentrator
== $ZZWAN.#S01
ALLOW 20 ERRORS
ABORT LINE $X25*
```

Printer Line Shutdown File

This example shows an SCF command file that stops the printer line associated with the SWAN concentrator \$ZZWAN.#S01 (configuration track-ID X001XX).

This file can be invoked automatically from the STOPSYS file, or you can invoke it by using the following TACL command:

```
> SCF/ IN $SYSTEM.SHUTDOWN.SDNLP, OUT $ZHOME /
== This is $SYSTEM.SHUTDOWN.SDNLP
== Shuts down the printer associated with the SWAN concentrator
== $ZZWAN.#S01
ALLOW 20 ERRORS
ABORT LINE $LP5516
```

Expand-Over-IP Line Shutdown File

This example shows an SCF command file that stops the Expand-over-IP communications line from \Case1, a NonStop S7000 server, to \Case2, a NonStop K-series server.

This file can be invoked automatically from the STOPSYS file, or you can invoke it by using the following TACL command:

```
> SCF/ IN $SYSTEM.SHUTDOWN.IP2CASE2, OUT $ZHOME /
== This is $SYSTEM.SHUTDOWN.IP2CASE2
ABORT LINE $Case2IP
```

Direct-Connect Line Shutdown File

This example shows an SCF command file that stops the direct-connect line on a SWAN concentrator.

This file can be invoked automatically from the STOPSYS file, or you can invoke it by using the following TACL command:

```
> SCF/ IN $SYSTEM.SHUTDOWN.STOPLH, OUT $ZHOME /
== This is $SYSTEM.SHUTDOWN.STOPLH
== This shuts down the direct-connect line
ALLOW 20 ERRORS
ABORT LINE $Case2elh
```

Spooler Shutdown File

This example shows a TACL command file that drains the spooler.

This file can be invoked automatically from the STOPSYS file, or you can invoke it by using the following TACL command:

```
> OBEY $SYSTEM.SHUTDOWN.SPLDRAIN
```

To maintain the integrity of the spooler environment, HP recommends that you wait until the spooler has finished draining rather than stop any spooler processes by using the TACL STOP command.

```
comment -- This is $SYSTEM.SHUTDOWN.SPLDRAIN
comment -- This file drains the spooler subsystem leaving all jobs intact.
SPOOLCOM $SPLS; SPOOLER, DRAIN
```

TMF Shutdown File

This example shows a TMFCOM command file that stops the Transaction Management Facility (TMF) subsystem.

This file can be invoked automatically from the STOPSYS file, or you can invoke it by using the following TACL command:

```
> TMFCOM / IN $SYSTEM.SHUTDOWN.TMFSTOP, OUT $ZHOME /
```

To maintain the integrity of the TMF environment, HP recommends that you wait until all transactions have finished rather than stop any TMF processes by using the TACL STOP command.

```
! comment -- This is $SYSTEM.SHUTDOWN.TMFSTOP
! comment -- This file stops any new transactions from being started,
! comment -- allows any transactions in process to finish, and then
! comment -- stops the TMF subsystem.

DISABLE BEGINTRANS; STOP TMF, WAIT ON; EXIT
```


17 Preventive Maintenance

[When to Use This Section](#) on page 17-1

[Monitoring Physical Facilities](#) on page 17-1

[Checking Air Temperature and Humidity](#) on page 17-1

[Checking Physical Security](#) on page 17-2

[Maintaining Order and Cleanliness](#) on page 17-2

[Checking Fire-Protection Systems](#) on page 17-2

[Cleaning System Components](#) on page 17-2

[Cleaning an Enclosure](#) on page 17-2

[Cleaning and Maintaining Printers](#) on page 17-2

[Cleaning Tape Drives](#) on page 17-3

[Handling and Storing Cartridge Tapes](#) on page 17-3

When to Use This Section

This section describes routine maintenance tasks required for Integrity NonStop NS-series servers.

Monitoring Physical Facilities

This subsection explains how to check the physical environment of your computer facility. You might be asked to monitor these aspects of your physical facility:

- Air temperature and humidity
- Physical security
- Order and cleanliness
- Fire-protection systems

Checking Air Temperature and Humidity

Check that the temperature and humidity are at the correct level established by management personnel. Monitor any sensors that control temperature and humidity. Your computer environment should have information posted that lists the names and telephone numbers of individuals that operators can call when a malfunction occurs with the heating, air conditioning, and humidity systems.

Checking Physical Security

Your company's security policy will guide you in the kind of security monitoring you perform. You might be asked to check doors and windows at the beginning and end of your shift and report the presence of unauthorized persons. In some facilities, operations staff might be responsible for monitoring and maintaining electronic security systems.

Maintaining Order and Cleanliness

Clutter and debris can cause accidents and fires. Dust, smoke, and spilled liquids can damage system hardware components. Depending on your company's policies, you might be asked to keep the computer room clean; inspect air filters; keep printer dust under control through periodic vacuuming; and enforce a ban on smoking, eating, and drinking in the computer room.

Checking Fire-Protection Systems

You might also be asked to check the fire alarms and fire extinguisher systems in your facility.

Cleaning System Components

This subsection contains basic information about cleaning enclosures, printers, and tape drives. Many companies have service-level agreements with HP that include regular preventive maintenance (PM) of their hardware components. If a Field Service Organization (FSO) representative handles cleaning and other preventive maintenance for your company, you need not be concerned with the cleaning tasks mentioned here.

Cleaning an Enclosure

Cleaning an enclosure is an infrequent task that you perform as required by conditions at your site. Many installations require only occasional dusting. To dust an enclosure, use a lint-free, antistatic dust cloth.

△ **Caution.** Do not use solvents or spray products on any part of an enclosure.

If you need to clean an enclosure, use a cotton cloth and a cleaning product formulated for computer equipment. Or use a damp cotton cloth and a mild, nonabrasive soap.

Cleaning and Maintaining Printers

Inspect all printers and replace the ribbons on line printers as needed. Replace the toner cartridges of laser printers that are shared by the user community when print quality lessens. To remove paper dust that can affect printer operation, vacuum printers periodically.

Cleaning Tape Drives

Clean tape drive heads and sensors frequently. For detailed information on cleaning tape drives, refer to the documentation shipped with your tape drive. How often you clean a tape drive or the tape path depends on use, operating environment, and tape quality. Cleaning supplies are available from HP. Use these materials:

- Cleaning solvent: HP supports the use of **only** isopropyl alcohol (91 percent or greater) as a tape-path cleaning solvent. Isopropyl alcohol cuts oil and grease, evaporates quickly, leaves no residue, and does not damage the tape path.
- Nonabrasive, lint-free cloths and swabs.
- A cleaning cartridge, which provides a safe, convenient way to clean some types of tape drives.

For ordering information, see the operator's guide shipped with the tape subsystem.

△ **Caution.** These precautions are extremely important to prevent damage:

- **Do not use cleaner solutions that contain lubricants.** Lubricants deposit a film on the tape head and impair performance.
 - **Do not use aerosol cleaners, even if they contain isopropyl alcohol.** The spray is difficult to control and often contains metallic particles that damage the tape head.
 - **Do not use soap and water on a tape path.** Soap leaves a thick film, and water can damage electronic parts.
 - **Do not use facial tissues.** Facial tissues leave highly abrasive lint on the tape path.
 - **Do not dip your cloths and swabs into the solvent can.** These items contaminate the solvent.
 - **Discard the cloths and swabs after use.** Even if they appear clean, cloths and swabs are contaminated by use. Dispose of these materials in an approved container.
-

Handling and Storing Cartridge Tapes

Cartridge tapes are nested in a cartridge by a spring-loaded mechanism that exposes the tape only when the cartridge tape is loaded into the tape drive. The drive mechanism uses a leader block to wind the tape through the tape path to load data. Exposing the tape or handling the cartridge tapes improperly can result in loss of data.

When handling and storing cartridge tapes:

- Protect cartridge tapes from shock or vibration.
- Do not expose cartridge tapes to high temperatures by leaving them in a car or in direct sunlight.
- Do not store cartridge tapes near magnetic fields such as those generated by power cables or computer monitors.

- Do not remove the leader block, pull out the tape, or press the reel lock. If the leader block is detached from the tape, contact the tape supplier for a leader block repair kit.
- When transporting cartridge tapes, do not stack the cartridges more than six high. Pack them carefully with the reel sides upright. The leader block edges can crack if they engage with each other.
- To store or transport tape cartridges in an ACL cartridge magazine, follow the same guidelines as you would for storing or transporting individual cartridge tapes.

A

Operational Differences Between Systems Running G-Series and H-Series RVUs

Users familiar with systems running G-series RVUs will find several major differences in the operational environment of systems running H-series RVUs. Although many of the operations to be performed remain the same, the tools you use to execute these operations might differ significantly. For H-series RVUs, these changes have been made:

- TSM is not supported in H-series. You must use OSM. Also, OSM's graphical representation of modular systems has a different look in H-series.
- In power failures, there is no memory hold-up for H-series. Ride-through is available only if the customer has a site uninterruptible power supply (UPS) or an in-cabinet UPS for all the affected cabinets.
- TAPEBOOT is not supported in H-series.
- In H-series, native compilers and linkers have new names. Therefore, automated scripts might require changing.
- Subvol for public libraries is `SYSnn` in G-series. In H-series, it is `ZDLLnnn` and requires changing scripts.
- REPLACEBOOT only applies to TNS and TNS/R. It does not apply to TNS/E.
- On G-series servers, the OSS shell command `ls` displays the contents of directories without visually distinguishing between subdirectories and files. On H-series servers, `ls` displays the contents of directories with a visual distinction between subdirectories and files -- subdirectory names are suffixed with a slash (/). This difference affects any OSS shell script that relied upon processing the output of the `ls` command.
- For H-series, DSM/SCM installation default is Manage OSS Files. For G-series, the default is not to manage OSS files.
- KMSF swap files have a larger memory size. It is now four times memory size per processor.
- Changes to automated debugging and dump mechanisms are required in H-series because of the new debuggers and debugger commands.
- The H-series OSS environment does not support TNS execution. OSS programs must be migrated to TNS/E native mode to run on an H-series system.

B

Tools and Utilities for Operations

[When to Use This Appendix](#) on page B-1

[BACKCOPY](#) on page B-2

[BACKUP](#) on page B-2

[Disk Compression Program \(DCOM\)](#) on page B-2

[Disk Space Analysis Program \(DSAP\)](#) on page B-2

[EMSDIST](#) on page B-2

[Event Management Service Analyzer \(EMSA\)](#) on page B-2

[File Utility Program \(FUP\)](#) on page B-3

[Measure](#) on page B-3

[MEDIACOM](#) on page B-3

[NonStop NET/MASTER](#) on page B-3

[NSKCOM and the Kernel-Managed Swap Facility \(KMSF\)](#) on page B-3

[OSM Package](#) on page B-3

[PATHCOM](#) on page B-4

[PEEK](#) on page B-4

[RESTORE](#) on page B-4

[SPOOLCOM](#) on page B-4

[Subsystem Control Facility \(SCF\)](#) on page B-4

[HP Tandem Advanced Command Language \(TACL\)](#) on page B-5

[TMFCOM](#) on page B-5

[Web ViewPoint](#) on page B-5

[ViewPoint](#) on page B-5

[ViewSys](#) on page B-6

When to Use This Appendix

This appendix briefly describes the tools and utilities that might be available on your system to assist you in performing the operations tasks for an Integrity NonStop NS-series server. The use of some of these tools and utilities is discussed throughout this guide. For a list of other documentation that provides detailed information about these tools and utilities, refer to [Appendix C, Related Reading](#).

BACKCOPY

Use the BACKCOPY utility to create one or two duplicate tapes for archive storage, distribution, or disaster recovery. You can also create one or two labeled (or unlabeled) tape sets from a labeled or unlabeled tape set. The BACKCOPY utility duplicates tapes that are made from a BACKUP utility file-mode operation, but it cannot duplicate tapes that are made from a BACKUP utility volume-mode operation.

BACKUP

Use the BACKUP utility to copy files from disk to magnetic tape.

Disk Compression Program (DCOM)

The Disk Compression Program (DCOM) moves disk-file extents to yield more usable space on a disk. Use the DCOM utility to analyze the current space allocation on a disk, relocate file extents on a disk, and reduce the number of free-space extents. You can also combine free space into larger extents so that files can be allocated with larger extents, which decreases the incidence of file-system error 43 (unable to obtain disk space for file extent).

Disk Space Analysis Program (DSAP)

The Disk Space Analysis Program (DSAP) analyzes how disk space is used on a specified volume. The DSAP utility copies the disk directory and free-space table to the current work file. By specifying options, you can manipulate this data to produce several different reports about the use of the disk space for that volume. The free-space table is limited only by your primary (main) and secondary (contiguous disk space) memory requirements.

EMSDIST

The EMSDIST program is the object program for a printing, forwarding, or consumer distributor, any of which you can start with a TACL RUN command.

Event Management Service Analyzer (EMSA)

Use the Event Management Service Analyzer (EMSA) to extract specific types of event messages from EMS log files and to create an Enscribe database that you can query to analyze problem trends.

File Utility Program (FUP)

The File Utility Program (FUP) is a component of the standard software package for the NonStop Kernel operating system. FUP software is designed to help you manage disk files, nondisk devices (printers, terminals, and tape drives), and processes (running programs) on the Integrity NonStop system. You can use FUP to create, display, and duplicate files; load data into files; alter file characteristics; and purge files.

Measure

Use the Measure program to collect and display system performance statistics about processors, processes, communication and network lines, files, disks, and terminals. Operations management personnel often use Measure to help fine-tune and balance a system.

MEDIACOM

MEDIACOM is the operator interface to the Distributed Systems Management/Tape Catalog (DSM/TC). It allows you to perform routine tape and tape-drive management operations.

NonStop NET/MASTER

NonStop NET/MASTER is used to integrate system and network management services. It serves as an alternative to the ViewPoint console application.

NSKCOM and the Kernel-Managed Swap Facility (KMSF)

NSKCOM is the command interface to the Kernel-Managed Swap Facility (KMSF). NSKCOM allows you to configure and manage permanently allocated swap files.

OSM Package

The HP Open System Management (OSM) product is the system management tool for Integrity NonStop systems. OSM offers a browser-based interface that improves scalability and performance and overcomes other limitations that existed in TSM. The OSM Low-Level Link has a new System Load dialog specifically for Integrity NonStop servers.

TSM does not supported Integrity NonStop NS-series servers; OSM is required.

For more information on the OSM package, including a description of the individual applications and how they differ from their TSM counterparts, see the *OSM Migration and Configuration Guide* and the *OSM Service Connection User's Guide*.

PATHCOM

PATHCOM is the interactive interface to the PATHMON process, through which users enter commands to configure and manage Pathway applications.

PEEK

Use the PEEK program to gather statistical information about processor activity, system storage pools, paging activity, message information, send instructions, and interrupt conditions.

RESTORE

Use the RESTORE utility to copy files from magnetic tape to disk.

SPOOLCOM

SPOOLCOM allows you to perform these tasks related to printing:

- Display the status of collectors, devices, print jobs, print processes, routing structures, and the spooler itself
- Change the location, state, or any attribute of your job
- Delete your print job from the spooler subsystem
- Restart a device that has gone offline with a device error

Subsystem Control Facility (SCF)

SCF configures and manages several subsystems that control system processes and hardware, including communications paths, disks, tapes, terminals, printers, and communications lines. You can run SCF from any workstation or terminal on the system after you are logged on. Use SCF to:

- Configure and add an object
- Remove an object
- Begin or restore access to an object
- Stop access to an object
- Show static configuration information for an object
- Show dynamic information for an object
- Automate subsystem startup and shutdown procedures
- Power off the system

HP Tandem Advanced Command Language (TACL)

The TACL product is the command interface to the NonStop Kernel operating system. In addition to providing full command-interpreter facilities, you can program the TACL interface to help you manage your system in these ways:

- Automate subsystem startup and shutdown procedures. For example, you can use TACL statements to initialize Pathway, the TMF subsystem, the TRANSFER system, and other subsystems.
- Run utilities and issue commands with either a fixed set of commands or a flexible set that you can tailor at run time.
- Create a customized environment that simplifies commonly performed tasks for users.

TMFCOM

TMFCOM allows you to enter commands that initiate communication with TMF, request various TMF operations, and terminate communication with TMF.

Web ViewPoint

Use Web ViewPoint, a browser-based product, to access the Event Viewer, Object Manager, and Performance Monitor subsystems. Web ViewPoint monitors and displays EMS events; identifies and lists all supported subsystems; manages NonStop server subsystems and user applications in a secure, automated, and customizable way; monitors and graphs performance attributes and trends; investigates and displays most active system processes; and offers simple navigation and a point-and-click command interface.

ViewPoint

Use ViewPoint to display event messages about current or past events occurring anywhere in the network on a set of block-mode events screens. The messages can be errors, failures, warnings, and requests for operator actions. The events screens allow operators to monitor significant occurrences or problems in the network as they occur. Critical events or events requiring immediate action are highlighted.

ViewSys

ViewSys is a system resource monitor that displays processor performance statistics and resource consumption for a set polling period. It updates the numbers automatically at the end of each polling period, which allows you to evaluate the effects of changes as those changes are made. ViewSys indicates the current allocation of a given resource and the percentage of that resource used. Thus, possible resource contention problems can be detected before they become serious.

Viewing the resource allocations across processors on a running system allows you to balance the application load more evenly. It can help you decide when to move user processes to processors and disk files that are less busy or when to relocate partitions to disk volumes that are less busy.

C Related Reading

For more information about tools and utilities used for system operations, refer to the documentation listed in [Table C-1](#).

Table C-1. Related Reading for Tools and Utilities (page 1 of 5)

Tool	Documentation	Description
BACKCOPY BACKUP DCOM DSAP	<i>Guardian Disk and Tape Utilities Manual</i>	This manual describes these disk and tape utilities: BACKCOPY, BACKUP, DCOM, DSAP, and RESTORE. This manual supports D-Series, G-series, and H-series RVUs.
EMSA	<i>Event Management Service (EMS) Analyzer Manual</i>	This manual describes how to specify parameters, such as subsystem ID, event number, text, start time and stop time, through the EMS conversational interface. The EMS Analyzer selects events from EMS log files.
EMSDIST	<i>Guardian User's Guide</i>	This guide describes how to use EMSDIST to display operator messages with a printing distributor, direct messages to a disk file, and print messages.
Measure	<i>Measure User's Guide</i> <i>Measure Reference Manual</i>	This manual describes how to use the Measure performance monitor to collect and examine system performance data. This manual describes the commands, callable procedures, and error messages of the Measure performance monitor.
MEDIACOM	<i>DSM/Tape Catalog User's Guide</i> <i>DSM/Tape Catalog Operator Interface (MEDIACOM) Manual</i> <i>Guardian User's Guide</i>	This guide describes the Distributed Systems Management/Tape Catalog (DSM/TC) software product, which allows users to organize, manage, and track tape volumes. It describes the components of DSM/TC and provides instructions and examples of how to configure, run, and maintain the DSM/TC system. This manual explains how to run a MEDIACOM session and describes the purpose and the syntax of the MEDIACOM commands. This guide contains information explaining how to perform routine operations relating to the tapes and tape drives on your system. The guide explains the MEDIACOM utility and provides examples for using it.

Table C-1. Related Reading for Tools and Utilities (page 2 of 5)

Tool	Documentation	Description
NonStop NET/MASTER	<i>NonStop NET/MASTER MS System Management Guide</i>	Describes how to integrate system and network management services. It serves as an alternative to the ViewPoint console application.
NSKCOM	<i>Kernel-Managed Swap Facility (KMSF) Manual</i>	This manual describes the operation of and command syntax for NSKCOM, the command interface to KMSF.
OSM package	<p><i>OSM Service Connection User's Guide</i> (also available as online help within OSM Service Connection)</p> <p><i>OSM Migration and Configuration Guide</i></p> <p><i>NonStop System Console Installer Guide</i></p> <p>Online help</p>	<p>This guide includes:</p> <ul style="list-style-type: none"> ● An overview of all OSM applications and components ● How to use the OSM Service Connection (the primary OSM interface) to monitor and perform actions on system and cluster resources <p>This guide includes:</p> <ul style="list-style-type: none"> ● Comparison of OSM and TSM software ● Hardware for which OSM is required ● System console hardware and software requirements for using OSM ● Coexistence and fallback issues ● How to migrate an existing TSM system list for OSM use <p>How to configure and start OSM server-side processes</p> <p>This guide describes how to install OSM client-based components and other required system console software.</p> <p>Online help is also available from within each of these OSM applications:</p> <ul style="list-style-type: none"> ● OSM Low-Level Link ● OSM Notification Director ● OSM Event Viewer ● Individual OSM guided procedures

Table C-1. Related Reading for Tools and Utilities (page 3 of 5)

Tool	Documentation	Description
PATHCOM	<i>TS/MP System Management Manual</i>	This manual describes the interactive management interface to the NonStop TS/MP product. It is intended for system managers and operators. It provides guidelines for configuring and controlling a NonStop TS/MP transaction processing system and its objects and for monitoring the status and performance of objects controlled by PATHMON in a Pathway environment. It also provides syntax for all relevant PATHCOM commands, as well as cause, effect, and recovery information for all PATHMON, PATHCOM, and LINKMON error messages.
PEEK	<i>PEEK Reference Manual</i>	This manual describes PEEK, a utility used to monitor statistical data about processors.
RESTORE	<i>Guardian Disk and Tape Utilities Reference Manual</i>	This manual describes these disk and tape utilities: BACKCOPY, BACKUP, DCOM, DSAP, and RESTORE. This manual supports both D-series, G-series, and H-series RVUs.
SPOOLCOM	<i>Spooler Plus Utilities Reference Manual</i> <i>Guardian User's Guide</i>	<p>This manual describes the spooler utilities—Peruse, SPOOLCOM, Font, and RPSetup— and presents the complete syntax for these utilities. It also presents a general introduction to the Spooler Plus subsystem.</p> <p>This guide contains information explaining how to perform routine spooler operations. It provides background information on spooler components and tells you how to use SPOOLCOM to monitor and manage your system's spooler operations. It includes guidelines for identifying and solving some common problems that can occur with your spooler subsystem and the printers associated with it.</p>
Startup and shutdown command files	<i>Integrity NonStop NS-Series Planning Guide</i>	This guide describes how to automate startup and shutdown procedures.

Table C-1. Related Reading for Tools and Utilities (page 4 of 5)

Tool	Documentation	Description
Subsystem Control Facility (SCF)	<i>SCF Reference Manual for H-Series RVUs</i>	This manual describes the operation of SCF on H-series RVUs and how it is used to configure, control, and monitor subsystems supported by an SCF interface.
SCF interface to the Kernel subsystem	<i>SCF Reference Manual for the Kernel Subsystem</i>	This manual describes the Kernel subsystem, and the configuration and management tasks you can perform using the SCF interface to the Kernel subsystem.
SCF interface to the storage subsystem	<i>SCF Reference Manual for the Storage Subsystem</i>	This manual describes how to use SCF to configure, control, and monitor storage devices.
SCF interface to the SLSA subsystem	<i>LAN Configuration and Management Manual</i>	This manual describes how to configure, operate, and manage the ServerNet LAN Systems Access (SLSA) subsystem. This manual includes detailed descriptions of the SCF commands used with the SLSA subsystem and a quick-reference section showing SCF command syntax.
SCF interface to the WAN subsystem	<i>WAN Subsystem Configuration and Management Manual</i>	This manual describes how to configure a ServerNet wide area network (SWAN) concentrator on an Integrity NonStop server. It also describes how to monitor, modify, and control the WAN subsystem. It includes detailed descriptions of the SCF commands used with the WAN subsystem.
SCF interface to other subsystems	Titles vary	These documents describe how to use the SCF interface to other subsystems.
TACL	<i>TACL Reference Manual</i>	This manual provides information on using the TACL interface.
TMFCOM	<i>TMF Operations and Recovery Guide</i>	This manual describes how to operate TMF and recover from error conditions. It is intended for those responsible for TMF system maintenance.
	<i>TMF Reference Manual</i>	This manual describes how to use the TMFCOM command interface to TMF. This manual includes command syntax, semantics, and examples, and is intended for system managers and operators.

Table C-1. Related Reading for Tools and Utilities (page 5 of 5)

Tool	Documentation	Description
Web ViewPoint	<i>Web ViewPoint Version 5.2 User Guide</i>	<p>This guide describes Web ViewPoint, a browser-based automated operation and management product that provides access to the Event Viewer, Object Manager, and Performance Monitor subsystems. Key features:</p> <ul style="list-style-type: none"> ● Monitors and displays EMS events ● Identifies and lists all supported subsystems ● Manages NonStop server subsystems and user applications in a secure, automated, and customizable way ● Monitors and graphs performance attributes and trends ● Investigates and displays most active system processes ● Offers simple navigation and a point-and-click command interface
ViewPoint	<i>ViewPoint Manual</i>	<p>This manual describes ViewPoint, a multifunction operations console application that allows the management of a network of systems. The manual contains information on installing, configuring, and starting ViewPoint for custom applications. It also describes the concepts underlying ViewPoint operation.</p> <p>Although the ViewPoint Manual applies to D-series, G-series, and H-series RVUs, some information might apply only to D- and G-series RVUs.</p>
ViewSys	<i>ViewSys User's Guide</i>	<p>This guide describes the operation of ViewSys and interpretation of the program output.</p>

D Converting Numbers

[When to Use This Appendix](#) D-1

[Overview of Numbering Systems](#) D-2

[Binary to Decimal](#) D-3

[Octal to Decimal](#) D-4

[Hexadecimal to Decimal](#) D-5

[Decimal to Binary](#) D-7

[Decimal to Octal](#) D-8

[Decimal to Hexadecimal](#) D-9

When to Use This Appendix

Refer to this appendix if you need to convert numbers from one numbering system to another.

Overview of Numbering Systems

Internally, a computer stores data as a series of off and on values represented symbolically by the binary digits, or bits, 0 and 1, respectively. Because numbers represented as strings of binary 0s and 1s are difficult to read, binary numbers are generally converted into octal, decimal, or hexadecimal form. [Table D-1](#) describes the binary, octal, decimal, and hexadecimal number systems.

Table D-1. Descriptions of Number Systems

Number System	Base	Description
Binary	2	Binary numbers are made up of the digits 0 and 1.
Octal	8	Octal numbers are made up of the digits 0, 1, 2, 3, 4, 5, 6, and 7.
Decimal	10	Decimal numbers are made up of the digits 0, 1, 2, 3, 4, 5, 6, 7, 8, and 9
Hexadecimal	16	Hexadecimal numbers are made up of the digits 0, 1, 2, 3, 4, 5, 6, 7, 8, and 9 and the letters A, B, C, D, E, and F.

In manuals for the NonStop server, a percent sign precedes a number that is not in decimal form:

- The % notation precedes an octal number.
- The %B notation precedes a binary number.
- The %H notation precedes a hexadecimal number. On some system displays, hexadecimal numbers are preceded by the notation 0X instead of %H.

Binary to Decimal

To convert a binary number to a decimal number:

1. Starting from the right, multiply the least significant (rightmost) binary digit by the first placeholder value. Moving towards the left, multiply each new binary digit by its corresponding placeholder value until the binary number is exhausted.

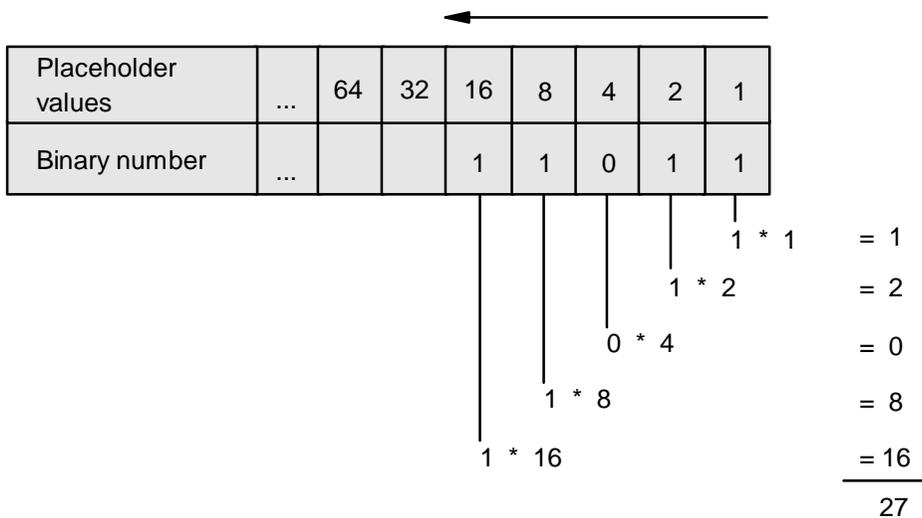
To establish placeholder values, the first placeholder value (on the far right) is 1. Then for each new placeholder value to the left, multiply the value to the right by 2.

2. Add the results of the multiplications in Step 1.

Example

Convert the binary value 11011 to its decimal equivalent. (In this example, the symbol “*” indicates multiplication.) Refer to [Figure D-1](#).

Figure D-1. Binary to Decimal Conversion



CDT 607.CDD

1. Take the rightmost binary digit and multiply it by the rightmost placeholder value.
2. Moving to the left, take the next binary digit and multiply it by the next placeholder value. Continue to do this until the binary number has been exhausted.
3. Add the multiplied values together. The result is:

Binary Value	Decimal Value
%B11011	27

Octal to Decimal

To convert an octal number to a decimal number:

1. Starting from the right, multiply the least significant (rightmost) octal digit by the first placeholder value. Moving towards the left, multiply each new octal digit by its corresponding placeholder value until the octal number is exhausted.

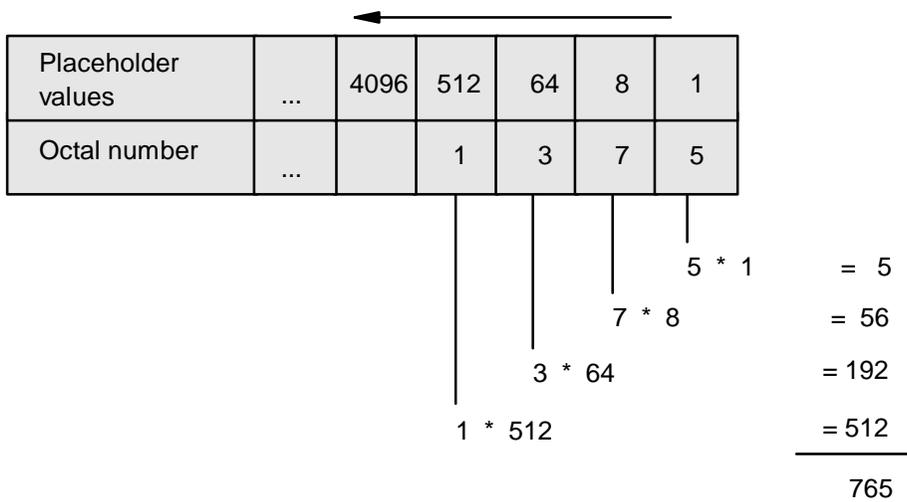
To establish placeholder values, the first placeholder value on the far right is 1. Then for each new placeholder value to the left, multiply the value to the right by 8.

2. Add the results of the multiplications in Step 1.

Example

Convert the octal value 1375 to its decimal equivalent. (In this example, the symbol “*” indicates multiplication.) Refer to [Figure D-2](#).

Figure D-2. Octal to Decimal Conversion



CDT 608CDD

1. Take the rightmost octal digit and multiply it by the rightmost placeholder value.
2. Moving to the left, take the next octal digit and multiply it by the next placeholder value. Continue to do this until the octal number has been exhausted.
3. Add the multiplied values together. The result is:

Octal Value	Decimal Value
%1375	765

Hexadecimal to Decimal

To convert a hexadecimal number to a decimal number:

1. Starting from the right, multiply the least significant (rightmost) hexadecimal digit by the first placeholder value. Moving towards the left, multiply each new hexadecimal digit by its corresponding placeholder value until the hexadecimal number is exhausted.

To establish placeholder values, the first placeholder value (on the far right) is 1. Then for each new placeholder value to the left, multiply the value to the right by 16.

Convert the letters of a hexadecimal number to decimal values before multiplying. Use this table for conversion:

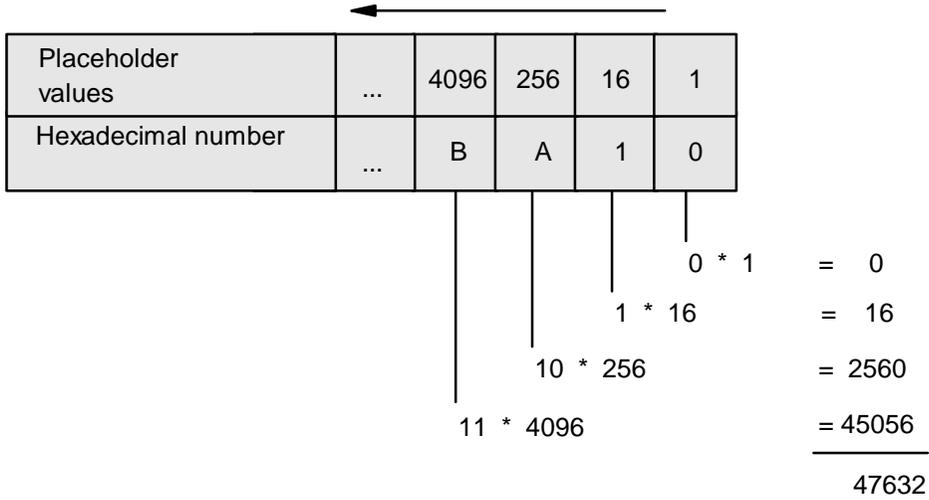
Hexadecimal	Decimal
A	10
B	11
C	12
D	13
E	14
F	15

2. Add the results of the multiplications in Step 1.

Example

Convert the hexadecimal value BA10 to its decimal equivalent. (In this example, the symbol “*” indicates multiplication.) Refer to [Figure D-3](#).

Figure D-3. Hexadecimal to Decimal Conversion



CDT 609.CDD

1. Take the rightmost hexadecimal digit and multiply it by the rightmost placeholder value.
2. Moving to the left, take the next hexadecimal digit and multiply it by the next placeholder value. Continue to do this until the hexadecimal number has been exhausted. Convert the hexadecimal digits A and B to their decimal values 10 and 11 before multiplying.
3. Add the multiplied values together. The result is:

Hexadecimal Value	Decimal Value
%HBA10	47632

Decimal to Binary

To convert a decimal number to a binary number:

1. Divide the decimal number by 2. The remainder of this first division becomes the least significant (rightmost) digit of the binary value.
2. Divide the quotient from Step 1 by 2, and use the remainder of the next division as the next digit (to the left) of the binary value. Continue to divide the quotients by 2 until the decimal number is exhausted. The remainder from the last division is the most significant (leftmost) digit of the binary value.

Example

Convert the decimal value 354 to its binary equivalent. (In this example, the symbol “/” indicates division.)

Step	Division	Quotient	Remainder	
1.	354/2	= 177	0	remainder = least significant (rightmost) digit
2.	177/2	= 88	1	
3.	88/2	= 44	0	
4.	44/2	= 22	0	
5.	22/2	= 11	0	
6.	11/2	= 5	1	
7.	5/2	= 2	1	
8.	2/2	= 1	0	
9.	1/2	= 0	1	remainder = most significant (leftmost) digit

The result is:

Decimal Value	Binary Value
354	%B101100010

Decimal to Octal

To convert a decimal number to an octal number:

1. Divide the decimal number by 8. The remainder of this first division becomes the least significant (rightmost) digit of the octal value.
2. Divide the quotient from Step 1 by 8, and use the remainder of the next division as the next digit (to the left) of the octal value. Continue to divide the quotients by 8 until the decimal number is exhausted. The remainder from the last division is the most significant (leftmost) digit of the octal value.

Example

Convert the decimal value 358 to its octal equivalent. (In this example, the symbol “/” indicates division.)

Step	Division	Quotient	Remainder	
1.	358/8	= 44	6	remainder = least significant (rightmost) digit
2.	44/8	= 5	4	
3.	5/8	= 0	5	remainder = most significant (leftmost) digit

The result is:

Decimal Value	Octal Value
358	%546

Decimal to Hexadecimal

To convert a decimal number to a hexadecimal number:

1. Divide the decimal number by 16. The remainder of this first division becomes the least significant (rightmost) digit of the hexadecimal value. If the remainder exceeds 9, convert the 2-digit remainder to its hexadecimal letter equivalent. Use this table for conversion.

Decimal	Hexadecimal
10	A
11	B
12	C
13	D
14	E
15	F

2. Divide the quotient from Step 1 by 16, and use the remainder of this next division as the next digit (to the left) of the hexadecimal value (converting 2-digit remainders as necessary). Continue to divide the quotients by 16 until the decimal number is exhausted. The remainder from the last division is the most significant (leftmost) digit of the hexadecimal value.

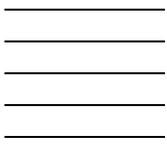
Example

Convert the decimal value 47632 to its hexadecimal equivalent. (In this example, the symbol “/” indicates division.)

Step	Division	Quotient	Remainder	
1.	47632/16	= 2977	0	remainder = least significant (rightmost) digit
2.	2977/16	= 186	1	
3.	186/16	= 11	10 = A	
4.	11/16	= 0	11 = B	remainder = most significant (leftmost) digit

The result is:

Decimal Value	Hexadecimal Value
47632	%HBA10



Safety and Compliance

This section contains three types of required safety and compliance statements:

- Regulatory compliance
- Waste Electrical and Electronic Equipment (WEEE)
- Safety

Regulatory Compliance Statements

The following 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.

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 Règlement 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 of 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.

European Union Notice

Products with the CE Marking comply with both the EMC Directive (89/336/EEC) and the Low Voltage Directive (73/23/EEC) issued by the Commission of the European Community.

Compliance with these directives implies conformity to the following European Norms (the equivalent international standards are in parenthesis):

- EN55022 (CISPR 22)—Electromagnetic Interference
- EN55024 (IEC61000-4-2, 3, 4, 5, 6, 8, 11)—Electromagnetic Immunity
- EN61000-3-2 (IEC61000-3-2)—Power Line Harmonics
- EN61000-3-3 (IEC61000-3-3)—Power Line Flicker
- EN60950-1 (IEC60950-1)—Product Safety

Laser Compliance

This product may be provided with an optical storage device (that is, CD or DVD drive) and/or fiber optic transceiver. Each of these devices contains a laser that is classified as a Class 1 Laser Product in accordance with US FDA regulations and the IEC 60825-1. The product does not emit hazardous laser radiation.



WARNING: Use the controls or adjustments or performance of procedures other than those specified herein or in the laser product's installation guide may result in hazardous radiation exposure. To reduce the risk of exposure to hazardous radiation:

- **Do not try to open the module enclosure. There are no user-serviceable components inside.**
- **Do not operate controls, make adjustments, or perform procedures to the laser device other than those specified herein.**
- **Allow only HP Authorized Service technicians to repair the module.**

The Center for Devices and Radiological Health (CDRH) of the U.S. Food and Drug Administration implemented regulations for laser products on August 2, 1976. These regulations apply to laser products manufactured from August 1, 1976. Compliance is mandatory for products marketed in the United States.

SAFETY CAUTION

The following icon or caution statements may be placed on equipment to indicate the presence of potentially hazardous conditions:



DUAL POWER CORDS CAUTION:

"THIS UNIT HAS MORE THAN ONE POWER SUPPLY CORD. DISCONNECT ALL POWER SUPPLY CORDS TO COMPLETELY REMOVE POWER FROM THIS UNIT."

"ATTENTION: CET APPAREIL COMPORTE PLUS D'UN CORDON D'ALIMENTATION. DÉBRANCHER TOUS LES CORDONS D'ALIMENTATION AFIN DE COUPER COMPLÈTEMENT L'ALIMENTATION DE CET ÉQUIPEMENT".

DIESES GERÄT HAT MEHR ALS EIN NETZKABEL. VOR DER WARTUNG BITTE ALLE NETZKABEL AUS DER STECKDOSE ZIEHEN.



Any surface or area of the equipment marked with these symbols indicates the presence of electric shock hazards. The enclosed area contains no operator-serviceable parts.

WARNING: To reduce the risk of injury from electric shock hazards, do not open this enclosure.

NOT FOR EXTERNAL USE

CAUTION: NOT FOR EXTERNAL USE. ALL RECEPTACLES ARE FOR INTERNAL USE ONLY.

ATTENTION: NE PAS UTILISER A L'EXTERIEUR DE L'EQUIPEMENT

IMPORTANT: TOUS LES RECIPIENTS SONT DESTINES UNIQUEMENT A UN USAGE INTERNE.

VORSICHT: ALLE STECKDOSEN DIENEN NUR DEM INTERNEN GEBRAUCH.

HIGH LEAKAGE CURRENT

To reduce the risk of electric shock due to high leakage currents, a reliable grounded (earthed) connection should be checked before servicing the power distribution unit (PDU).

Observe the following limits when connecting the product to AC power distribution devices: For PDUs that have attached AC power cords or are directly wired to the building power, the total combined leakage current should not exceed 5 percent of the rated input current for the device.

“HIGH LEAKAGE CURRENT, EARTH CONNECTION ESSENTIAL BEFORE CONNECTING SUPPLY”

“HOHER ABLEITSTROM. VOR INBETRIEBNAHME UNBEDINGT ERDUNGSVERBINDUNG HERSTELLEN”

“COURANT DE FUITE E'LEVE'. RACCORDEMENT A LA TERRE INDISPENSABLE AVANT LE RACCORDEMENT AU RESEAU”

FUSE REPLACEMENT

CAUTION – For continued protection against risk of fire, replace fuses only with fuses of the same type and the same rating. Disconnect power before changing fuses.

Waste Electrical and Electronic Equipment (WEEE)

Information about the Waste Electrical and Electronic Equipment (WEEE) directive is available from the NonStop Technical Library (NTL) home page. Select **Safety and Compliance > Waste Electrical and Electronic Equipment (WEEE)**.

Important Safety Information

Safety information is available from the NTL home page. Select **Safety and Compliance > Important Safety Information**. To open the safety information in a language other than English, select the language. Local HP support can also help direct you to your safety information.

Index

Numbers

4-Port ServerNet Extender (4PSE) [2-2](#),
[3-5](#), [7-3](#), [7-4](#), [8-3](#), [Index-1](#)

4-Port ServerNet Extender (4PSE),
servicing [8-7](#)

A

Asynchronous Terminal Process 6100
(ATP6100) [6-3](#)

ATM 3 ServerNet adapter (ATM3SA) [6-2](#)

ATM3SA [6-2](#)

ATP6100 [6-3](#), [16-21](#)

Automating

 system shutdown [16-3](#)

 system startup [16-2](#)

B

BACKCOPY utility [B-2](#)

BACKUP utility

 backing up configuration and operations
 files [9-20](#)

 description of [B-2](#)

Batteries

 charging [14-4](#)

 maintaining [14-4](#)

Battery ride-through [15-22](#)

Binary number system [D-2](#)

Binary to decimal conversion [D-3](#)

Bus dumps

 See Dumps

C

Cartridge tape, handling and storing [17-3](#)

CIIN file

 contents [16-5](#)

 establishing [16-6](#)

 file name [16-6](#)

 initial location of [16-2](#)

 modifying [16-6](#)

 ownership [16-6](#)

 security [16-6](#), [16-7](#)

 specifying [16-6](#)

 system behavior when absent [16-7](#)

Cleaning enclosures [17-2](#)

Collector (spooler), checking status of [12-2](#)

Command files

 examples [16-4/16-23](#)

Communications line [6-10](#)

Communications Process subsystem
(CP6100) [6-3](#)

Compaq TSM

 See TSM

Configuration

 files

 CONFTEXT

 See CONFTEXT file

 INITIAL_COMMAND_FILE [16-6](#)

 TCP/IP stacks [16-15](#)

CONFLIST file [9-20](#)

CONFTEXT file [9-20](#)

 INITIAL_COMINT_INFILE [16-6](#)

 INITIAL_COMMAND_FILE [16-6](#)

Converting numbers

 See Number conversion

CP6100 [6-3](#)

CPU n has been dumped to dumpfile
message [9-18](#)

D

- DCOM [10-15](#), [B-2](#)
- Decimal number system [D-2](#)
- Decimal to binary conversion [D-7](#)
- Decimal to hexadecimal conversion [D-9](#)
- Decimal to octal conversion [D-8](#)
- Direct-connect line
 - shutdown file [16-22](#)
 - startup file [16-18](#)
- Disk Compression Program (DCOM) [10-15](#), [B-2](#)
- Disk drives
 - common problems [10-11](#)
 - description of [10-2](#)
 - LEDs [3-20](#)
 - monitoring [10-4](#)
 - recovery operations for [10-12](#), [10-13](#)
- Disk Space Analysis Program (DSAP) [B-2](#)
- Distributed Systems Management/Tape Catalog (DSM/TC) [B-3](#)
- DSAP [B-2](#)
- DSM/TC [B-3](#)
- Dumps
 - completed message [9-18](#)
 - dump file
 - checking with FUP [9-18](#)
 - submitting to service provider [9-21](#)
 - processor to disk [9-19](#)

E

- E4SA [6-2](#)
- EMS Analyzer (EMSA) [B-2](#)
- EMS event messages, monitoring [4-1/4-2](#)
- EMSA [B-2](#)
- EMSDIST
 - description of [B-2](#)
 - using to monitor EMS event messages [4-2](#)
- EMSLOG file [9-20](#)

- Enclosures
 - cleaning [17-2](#)
- Enterprise Storage System
 - See ESS
- ESS [8-2](#)
- Ethernet 4 ServerNet adapter (E4SA) [6-2](#)
- Event Management Service (EMS) [4-1](#)
- Examples
 - checking file size [10-10](#)
 - checking status of PATHMON process [13-6](#)
 - checking status of TMF [13-4](#)
 - MEDIACOM STATUS TAPE DRIVE command [11-5](#)
 - Problem-Solving Worksheet [1-4](#)
 - SCF STATUS ADAPTER command [6-5](#)
 - SCF STATUS DISK command [10-6](#)
 - SCF STATUS LIF command [6-6](#)
 - SCF STATUS LINE command [6-11](#)
 - SCF STATUS PIF command [6-6](#)
 - SCF STATUS TAPE command [11-5](#)
 - start-of-shift checklist [3-3](#)
- Expand-over-IP
 - startup file [16-18](#)

F

- Fast Ethernet ServerNet adapter (FESA) [6-2](#)
- FCDM [2-2](#)
- FCSA [6-2](#)
 - overview [8-2](#)
 - problems with [8-4](#)
 - states [8-4](#)
- FESA [6-2](#)
- Fibre Channel disk module (FCDM) [2-2](#)
- Fibre Channel ServerNet adapter
 - See FCSA
- Fibre Channel ServerNet adapter (FCSA) [6-2](#)

File Utility Program (FUP)
 description of [B-3](#)
 INFO command [9-18](#), [10-9](#)

Freeze
 enabling or disabling on a
 processor [9-15](#)
 freeze code=%nn message [9-8](#)
 hardware error [9-7](#)

FUP
 See File Utility Program (FUP)

G

G4SA [6-2](#)
 monitoring [8-5](#)
 overview [8-2](#)
 states [8-6](#)

GESA [6-2](#)

Gigabit Ethernet 4-port adapter (G4SA) [6-2](#)

Gigabit Ethernet ServerNet adapter [6-2](#)

Guided procedures, OSM [1-12](#)

G-series [-xv](#)

H

Halting processors [9-10](#)
 See also Processor halts

Hang
 of processor [9-7](#)
 of system, recovery operations for [9-10](#)

Hexadecimal number system [D-2](#)

Hexadecimal to decimal conversion [D-5](#)

Home terminal, using \$ZHOME [16-4](#)

Hometerm
 See VHS

HP NonStop Open System Management
 (OSM)
 See OSM

HP NonStop Transaction Management
 Facility (TMF)
 See TMF

HP Tandem Advanced Command
 Language (TACL) [9-22](#)
 See TACL

I

INFO command (FUP) [9-18](#), [10-9](#)
 INITIAL_COMINT_INFILE [16-6](#)
 INITIAL_COMMAND_FILE [16-6](#)
 Integrity NonStop NS1000 system [2-3](#)
 Integrity NonStop NS14000 system [7-1](#),
[7-2](#), [7-3](#), [7-4](#)
 Integrity NonStop NS14000 system, with
 IOAM [2-3](#)
 Integrity NonStop NS14000 system, with
 VIO [2-3](#)
 Integrity NonStop NS16000 system [2-2](#)
 IOAM enclosure [8-2](#)
 I/O adapter module (IOAM) [8-2](#)

K

Kernel-Managed Swap Facility (KMSF) [B-3](#)
 KMSF [B-3](#)

L

LEDs
 status [15-2](#)
 LEDs, status [3-20](#)
 LIFs [6-2](#)
 Logical interfaces (LIFs) [6-2](#)

M

Measure program [B-3](#)
 MEDIACOM
 description of [B-3](#)
 interface [11-9](#), [B-3](#)
 STATUS TAPEDRIVE command [11-6](#)

Monitoring

- communications subsystems [6-13](#)
 - disk drives [10-4](#)
 - EMS event messages [4-1/4-2](#)
 - G4SA [8-5](#)
 - overview [3-1/3-22](#)
 - printers [12-1](#)
 - processes [5-1/5-6](#)
 - processors [9-4/9-8](#)
 - ServerNet fabrics [7-4/7-7](#)
 - tape drives [11-1/11-7](#)
 - terminals [12-1](#)
- MSP 0 or 1 [15-23](#)

N

- NonStop NET/MASTER application [B-3](#)
- NonStop TCP/IP [6-3](#)
- NonStop TCP/IPv6 [6-3](#)
- NonStop Virtual Hometerm Subsystem (VHS) [16-4](#)
- NSAA (NonStop advanced architecture) [2-2](#)
- NSKCOM [B-3](#)
- NSVA (NonStop value architecture) [2-2](#)
- Number conversion
 - binary to decimal [D-3](#)
 - decimal to binary [D-7](#)
 - decimal to hexadecimal [D-9](#)
 - decimal to octal [D-8](#)
 - hexadecimal to decimal [D-5](#)
 - octal to decimal [D-4](#)
 - overview [D-1](#)

O

- Octal number system [D-2](#)
- Octal to decimal conversion [D-4](#)
- OSM
 - CIIN file [16-7](#)
 - description of [B-3](#)
 - documentation [C-2](#)
 - guided procedures [1-12](#)
 - launching [1-11](#)
 - security [16-7](#)
 - using to monitor and resolve problems [3-7](#)
- OSM Event Viewer [4-2](#)
- Outages, planned [15-14](#)

P

- PAM [6-3](#)
- Parallel Library TCP/IP [2-9](#)
- PATHCOM [13-4](#)
- PATHMON processes [13-4](#)
- Pathway
 - commands [13-4](#)
 - processes [13-4](#)
 - transaction processing applications [13-1](#), [13-4](#)
- PEEK program [B-4](#)
- Physical interfaces (PIFs) [6-3](#)
- PIFs [6-3](#)
- Planned outages [15-14](#)
- Port Access Method (PAM) [6-3](#)
- POSTs
 - See Power-on self-tests (POSTs)

Power failure

- how external devices respond to [14-2](#)
- preparing for
 - maintaining batteries [14-4](#)
 - monitor batteries [14-4](#)
 - monitor power supplies [14-4](#)
 - ride-through time [14-3](#)
- recovery operations [14-4](#)
- response
 - ESS cabinets [14-3](#)
 - external devices [14-2](#)
 - NonStop NS-series cabinets [14-2](#)
 - NonStop S-series enclosures [14-2](#)
 - systems [14-2](#)

Powering off the system [15-17](#)

Powering on

- external system devices [15-3](#)

Power-on self-tests (POSTs)

- system power-on [15-2](#)

Printers

- monitoring [12-1](#)
- recovery operations for [12-2](#)

Problems, common

- disk drive [10-11](#)
- tape drive [11-7](#)

Processes

- generic [5-2](#)
- I/O [5-2](#)
- monitoring [5-3/5-6](#)
- recovery operations for [5-6](#)
- system [5-1](#)

Processor halts

- halt code = %nn message [9-8](#)
- recovery operations for [9-9](#)

Processors

- dumps
 - See Dumps
- freeze
 - See Freeze
- halt
 - See Processor halts
- halting processors [9-10](#)
- hang [9-7](#)
- monitoring [3-12](#)
- recovery operations for [9-9/9-21](#)

RRCVDUMP utility [9-18](#)

Recovery operations

- for disk drives [10-12](#), [10-13](#)
- for printers [12-2](#)
- for processor halt [9-9](#)
- for processors [9-7/9-20](#)
- for ServerNet fabrics [7-8](#)
- for system console [1-3](#)
- for tape drives [11-8](#)
- for \$SYSTEM [15-20](#)

Reloading single processor on running server [9-19](#)RESTORE utility [B-4](#)

S

SACs [6-2](#)

SCF [B-4](#)

commands

HELP [3-7](#)

LISTDEV [2-7](#)

STATUS ADAPTER [6-4](#)

STATUS DISK [10-5](#)

STATUS LIF [6-6](#)

STATUS LINE [6-10](#)

STATUS PIF [6-5](#)

STATUS SAC [6-5](#)

STATUS TAPE [11-5](#)

STATUS, examples of [3-13](#)

managing disks [11-9](#)

powering off the system [15-17](#)

storage device recovery [10-12](#), [11-8](#)

using to solve problems [3-7](#)

ServerNet addressable controllers (SACs) [6-2](#)

ServerNet fabrics

monitoring [7-4/7-7](#)

recovery operations for [7-8](#)

ServerNet switch board [8-2](#)

Setting system time [14-5](#)

Shutdown files

about [16-19/16-23](#)

ATP6100 lines [16-21](#)

automating [16-3](#)

CP6100 lines [16-21](#)

Expand-over-IP lines [16-22](#)

security [16-19](#)

sequence [16-3](#), [16-19](#)

spooler [16-23](#)

system shutdown file [16-20](#)

TMF [16-23](#)

X.25 lines [16-21](#)

SNAX/APN [6-3](#)

SPOOLCOM [B-4](#)

Spooler [16-14](#), [16-23](#)

Startup files

about [16-5/16-18](#)

ATP6100 lines [16-17](#)

automating [16-2](#)

CIIN [16-2](#)

configuration database [16-12](#)

CP6100 [16-17](#)

direct-connect [16-18](#)

Expand-over-IP [16-18](#)

invoking [16-2](#)

security [16-11](#)

sequence [16-11](#)

spooler warm start [16-14](#)

system startup file [16-12](#)

TCP/IP stacks [16-11](#)

TMF warm start [16-14](#)

X.25 lines [16-17](#)

States

FCSA [8-4](#)

G4SA [8-6](#)

Stopping the system [15-16](#), [15-17](#)

Storing cartridge tapes [17-3](#)

Subsystem Control Facility (SCF)

See SCF

Subsystems

displaying configuration of [2-9](#)

Kernel [2-10](#)

SLSA [2-12](#), [6-2](#)

storage [2-10](#)

TCP/IP [2-9](#)

WAN [2-13](#), [6-2](#)

SWAN concentrator [16-14](#)

System

performance [15-14](#)

powering off [15-17](#)

recording configuration of [2-4](#)

stopping [15-16](#), [15-17](#)

System console

TACL window [16-5](#)

System console, recovery operations for [1-3](#)

System time, setting [14-5](#)

S-series [-xv](#)

T

TACL [9-22](#), [16-5](#), [B-5](#)

Tape drives

common problems [11-7](#)

monitoring [11-2](#)

recovery operations for [11-8](#)

Tapes, handling and storing [17-3](#)

TCP/IP

configuration file [16-15](#)

startup file [16-14](#)

TCP/IPv6 [6-3](#)

Terminals

monitoring [12-1](#)

recovery operations for [12-1](#)

Time, system, setting [14-5](#)

TMF [16-14](#), [16-23](#)

states of [13-3](#)

STATUS command [13-1](#), [13-2](#)

TMFCOM command [13-1/13-3](#)

Token-Ring ServerNet adapter (TRSA) [6-2](#)

TRSA [6-2](#)

TSM

CIIN file [16-7](#)

launching [1-11](#)

security [16-7](#)

TSM workstation

See System console

V

ViewPoint

description of [B-5](#)

using to monitor EMS event messages [4-2](#)

ViewSys utility [9-7](#), [B-6](#)

VIO enclosure, description [2-3](#)

VIO enclosure, powering on [15-3](#)

Virtual Hometerm Subsystem (VHS) [16-4](#)

W

Web ViewPoint

using to access the Event Viewer [4-2](#)

X

X.25 lines [16-21](#)

Special Characters

\$SYSTEM, recovery operations for [15-20](#)

\$YMIOP.#CLCI [16-3](#), [16-5](#)

\$YMIOP.#CNLSL [16-3](#)

\$ZHOME [16-4](#)

