

HP Integrity NonStop Operations Guide for H-Series and J-Series RVUs

HP Part Number: 529869-023
Published: February 2014
Edition: J06.03 and subsequent J-series RVUs and H06.13 and subsequent H-series RVUs



© Copyright 2014 Hewlett-Packard Development Company, L.P.

Legal Notice

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Export of the information contained in this publication may require authorization from the U.S. Department of Commerce.

Microsoft, Windows, and Windows NT are U.S. registered trademarks of Microsoft Corporation.

Intel, Pentium, and Celeron are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java® is a registered trademark of Oracle and/or its affiliates. Motif, OSF/1,

Motif, OSF/1, UNIX, X/Open, and the "X" device are registered trademarks, and IT DialTone and The Open Group are trademarks of The Open Group in the U.S. and other countries.

Open Software Foundation, OSF, the OSF logo, OSF/1, OSF/Motif, and Motif are trademarks of the Open Software Foundation, Inc. OSF MAKES NO WARRANTY OF ANY KIND WITH REGARD TO THE OSF MATERIAL PROVIDED HEREIN, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. OSF shall not be liable for errors contained herein or for incidental consequential damages in connection with the furnishing, performance, or use of this material.

© 1990, 1991, 1992, 1993 Open Software Foundation, Inc. The OSF documentation and the OSF software to which it relates are derived in part from materials supplied by the following: © 1987, 1988, 1989 Carnegie-Mellon University. © 1989, 1990, 1991 Digital Equipment Corporation. © 1985, 1988, 1989, 1990 Encore Computer Corporation. © 1988 Free Software Foundation, Inc. © 1987, 1988, 1989, 1990, 1991 Hewlett-Packard Company. © 1985, 1987, 1988, 1989, 1990, 1991, 1992 International Business Machines Corporation. © 1988, 1989 Massachusetts Institute of Technology. © 1988, 1989, 1990 Mentat Inc. © 1988 Microsoft Corporation. © 1987, 1988, 1989, 1990, 1991, 1992 SecureWare, Inc. © 1990, 1991 Siemens Nixdorf Informationssysteme AG. © 1986, 1989, 1996, 1997 Sun Microsystems, Inc. © 1989, 1990, 1991 Transarc Corporation. OSF software and documentation are based in part on the Fourth Berkeley Software Distribution under license from The Regents of the University of California. OSF acknowledges the following individuals and institutions for their role in its development: Kenneth C.R.C. Arnold, Gregory S. Couch, Conrad C. Huang, Ed James, Symmetric Computer Systems, Robert Elz. © 1980, 1981, 1982, 1983, 1985, 1986, 1987, 1988, 1989 Regents of the University of California.

Contents

About This Document	13
Supported Release Version Updates (RVUs).....	13
Intended Audience.....	13
New and Changed Information in 529869-023 Edition.....	13
New and Changed Information in 529869-022 Edition.....	13
New and Changed Information in 529869-021 Edition.....	14
New and Changed Information in 529869-020 Edition.....	15
New and Changed Information in 529869-019 Edition.....	15
New and Changed Information in 529869-018 Edition.....	18
New and Changed Information in 529869-017 Edition.....	18
New and Changed Information in 529869-016 Edition.....	19
New and Changed Information in 529869-015 Edition.....	20
Document Organization.....	20
Notation Conventions.....	21
General Syntax Notation.....	21
Related Information.....	23
Publishing History.....	24
HP Encourages Your Comments.....	24
1 Introduction to NonStop Operations	26
When to Use This Chapter.....	26
Understanding the Operational Environment.....	26
What Are the Operator Tasks?.....	26
Starting the NonStop System and Loading the NonStop Operating System.....	27
Updating Firmware.....	27
Configuring and Controlling System Components.....	27
Monitoring System Hardware and Performance Subsystems and Performing Recovery Operations	
.....	27
Preparing for and Recovering from Power Failures.....	28
Stopping and Powering Off the System.....	28
Powering On and Starting the System.....	28
Creating Startup and Shutdown Files.....	28
Performing Preventive Maintenance.....	28
Operating Disk Drives and Tape Drives.....	28
Responding to Spooler Problems.....	28
Logging On to a NonStop System.....	28
System Consoles.....	28
Opening a TACL Window.....	29
Opening a TACL Window Directly From MR-Win6530	29
Opening a TACL Window From the OSM Low-Level Link	29
Overview of OSM Applications.....	29
Launching OSM Applications.....	30
Service Procedures.....	31
Support and Service Collection.....	31
2 Determining Your System Configuration	32
When to Use This Chapter.....	32
NonStop BladeSystems.....	32
NonStop BladeSystems NB50000c, NB50000c-cg, NB54000c, NB54000c-cg, NB56000c,	
and NB56000c-cg Modular Hardware Components.....	32
NonStop BladeSystem Options.....	33
Differences Between NonStop BladeSystems.....	34

NonStop BladeSystems NB50000c and NB50000c-cg.....	34
NonStop BladeSystems NB54000c and NB55000c-cg.....	34
NonStop BladeSystems NB56000c and NB56000c-cg.....	34
NonStop NS-Series Systems.....	34
NonStop NS-Series Modular Hardware Components.....	34
NonStop NS-Series System Options.....	36
Differences Between NonStop NS-Series Systems.....	36
NonStop NS2400 Series Systems.....	36
NonStop NS2300 Series Systems.....	36
NonStop NS2200 Series Systems.....	36
NonStop NS2100 Series Systems.....	36
NonStop NS2000 Series Systems.....	36
NonStop NS16000 Series Systems.....	37
NonStop NS14000 Series Systems.....	37
NonStop NS1000 and NS1200 Series Systems.....	38
Terms Used to Describe System Hardware Components.....	38
Device.....	38
System Resource or Object.....	38
Server and System.....	38
Blade.....	39
Recording Your System Configuration.....	39
Using SCF to Determine Your System Configuration.....	39
SCF System Naming Conventions.....	40
SCF Configuration Files.....	40
Using SCF to Display Subsystem Configuration Information.....	41
SCF LISTDEV: Listing the Devices on Your System.....	41
Displaying SCF Configuration Information for Subsystems	43
CIP Subsystem.....	43
TCP/IP Subsystem.....	44
Kernel Subsystem.....	44
Storage Subsystem.....	44
ServerNet LAN Systems Access (SLSA) Subsystem.....	46
WAN Subsystem.....	46
Additional Subsystems Controlled by SCF.....	47
Displaying Configuration Information—SCF Examples.....	48
CIP Subsystem Configuration.....	50
3 Overview of Monitoring and Recovery.....	51
When to Use This Chapter.....	51
HP Tools for Monitoring System Resources.....	52
What the HP Tools Monitor.....	52
Monitored NonStop Hardware Resources.....	52
Using HP SIM.....	55
Using the HP SIM Plug-Ins.....	57
NonStop Software Essentials.....	57
NonStop I/O Essentials.....	57
NonStop Cluster Essentials.....	58
NonStop Cluster Performance Essentials.....	58
Insight Control Power Management.....	58
HP Insight Remote Support Advanced.....	59
Using the OSM Service Connection.....	59
Recovery Operations for Problems Detected by OSM.....	59
Monitoring Problem Incident Reports.....	59
Generating Diagnostic Data.....	60
Using SCF.....	60

Determining Device States.....	60
SCF Object States.....	62
Using Onboard Administrator.....	63
Automating Routine System Monitoring Procedures.....	63
Using Automated HP Manageability Tools.....	64
Using TACL Macros, Routines, or Command Files.....	64
Using the Status LEDs to Monitor the System	68
Related Reading for NonStop Monitoring.....	70
4 Monitoring the Performance of NonStop Subsystems.....	71
Summary of NonStop Performance Objects Monitored.....	71
HP Manageability Tools That Monitor NonStop Performance.....	72
NonStop Cluster Performance Essentials.....	72
NonStop Availability Statistics and Performance (ASAP).....	72
Web ViewPoint.....	73
Web ViewPoint NonStop Storage Analyzer Plug-In.....	74
HP Operations Agent for NonStop (OVNM).....	75
HP Performance Agent for NonStop (OVNPM).....	75
5 Monitoring EMS Event Messages.....	77
When to Use This Chapter.....	77
What Is the Event Management Service (EMS)?.....	77
Tools for Monitoring EMS Event Messages.....	77
HP SIM Event Status Column and Event Tab.....	77
OSM Event Viewer	77
Web ViewPoint.....	78
EMSDIST.....	78
Related Reading for EMS Event Messages.....	78
6 Processes: Monitoring and Recovery.....	79
When to Use This Chapter.....	79
Types of Processes.....	79
System Processes.....	79
I/O Processes (IOPs).....	79
CIP Processes.....	80
Generic Processes.....	80
Monitoring Processes.....	80
Automated HP Tools That Monitor Processes.....	81
Monitoring System Processes.....	81
Monitoring IOPs.....	82
Monitoring CIP Processes.....	82
Monitoring Generic Processes.....	82
Monitoring the Status of \$ZZKRN.....	82
Monitoring the Status of All Generic Processes.....	82
Recovery Operations for Processes.....	84
Related Reading for Processes.....	84
7 Communications Subsystems: Monitoring and Recovery.....	85
When to Use This Chapter.....	85
Communications Subsystems.....	85
Local Area Networks (LANs) and Wide Area Networks (WANs).....	85
Cluster I/O Modules (CLIMs).....	87
Monitoring Communications Subsystems and Their Objects.....	88
Automated HP Tools That Monitor Communications Subsystems.....	88
Monitoring the SLSA Subsystem.....	88
Monitoring the Status of an Adapter and Its Components.....	89
Monitoring the WAN Subsystem.....	90

Monitoring Status for a SWAN Concentrator.....	91
Monitoring Status for a Data Communications Device.....	91
Monitoring WAN Processes.....	92
Monitoring CLIPs.....	92
Monitoring the NonStop TCP/IP Subsystem.....	93
Monitoring the NonStop TCP/IP Process.....	93
Monitoring NonStop TCP/IP Routes.....	93
Monitoring NonStop TCP/IP Subnets.....	94
Monitoring the CIP Subsystem.....	94
Monitoring CLIM Status.....	96
Monitoring Line-Handler Process Status.....	96
Examples.....	97
Tracing a Communications Line.....	98
Recovery Operations for Communications Subsystems	98
Related Reading for Communications Subsystems.....	99
8 ServerNet Resources: Monitoring and Recovery.....	100
When to Use This Chapter.....	100
ServerNet Connectivity.....	100
ServerNet Communications Network.....	101
Monitoring the Status of the ServerNet Fabrics.....	102
Monitoring the ServerNet Fabrics Using OSM.....	102
Monitoring the ServerNet Fabrics Using SCF.....	103
Normal ServerNet Fabric States	104
Identifying ServerNet Fabric Problems.....	104
Recovery Operations for the ServerNet Fabrics	105
Recovery Operations for a Down Disk Due to a Fabric Failure.....	105
Recovery Operations for a Down Path Between Processors.....	105
Recovery Operations for a Down Processor.....	105
Recovery Operations for a File-System Error.....	105
Related Reading for ServerNet Resources.....	105
9 I/O Adapters and Modules: Monitoring and Recovery.....	106
When to Use This Chapter.....	106
I/O Adapters and Modules.....	106
Fibre Channel ServerNet Adapter (FCSA).....	107
Gigabit Ethernet 4-Port Adapter (G4SA).....	107
4-Port ServerNet Extender (4PSE).....	108
Monitoring I/O Adapters and Modules.....	108
Monitoring the FCSAs.....	108
Monitoring the G4SAs.....	109
Monitoring the 4PSEs.....	110
Recovery Operations for I/O Adapters and Modules.....	110
Related Reading for I/O Adapters and Modules.....	111
10 CLuster I/O Modules (CLIMs): Monitoring and Recovery.....	112
When To Use This Section.....	112
CLuster I/O Modules (CLIMs).....	112
Monitoring the CLIMs with HP SIM.....	112
Monitoring the CLIMs with OSM.....	112
Monitoring the CLIMs with SCF.....	112
CLIM Device States.....	114
Using SCF to Monitor a CLIM with \$ZZCIP.....	115
Recovery Operations for CLIMs.....	118
Related Reading for CLIMs.....	119

11 Processors and Components: Monitoring and Recovery.....	120
When to Use This Chapter.....	120
Overview of Processors.....	120
Processors in NonStop Systems Running J-Series RVUs.....	121
NonStop NS16000 Series and NS14000 Series Blade Complex Processors.....	122
NonStop NS1200 and NS1000 Blade Element Processors.....	124
Monitoring and Maintaining Processors.....	124
Monitoring Processors Automatically Using TFDS.....	125
Monitoring Processor Status Using the OSM Low-Level Link.....	125
Monitoring Processor Status Using the OSM Service Connection.....	125
Managing and Setting NonStop BladeSystems NB54000c, NB54000c-cg, NB56000c, and NB56000c-cg Processor Power Usage with Insight Control Power Management.....	127
Monitoring Processor Performance Using ViewSys.....	128
Identifying Processor Problems.....	128
Processor or System Hangs.....	128
Processor Halts.....	128
OSM Alarms and Attribute Values.....	129
Recovery Operations for Processors.....	129
Recovery Operations for a Processor Halt.....	130
Halting One or More Processors.....	130
Reloading a Single Processor on a Running system.....	130
Using TACL RELOAD to Perform Reload.....	131
Using the OSM Service Connection to Perform Reload.....	132
Recovery Operations for a System Hang.....	133
Enabling/Disabling Processor and System Freeze.....	134
Freezing the System and Freeze-Enabled Processors.....	134
Dumping a Processor to Disk.....	134
Loading a Down System to Perform a Processor Dump to Disk.....	135
Using RCVDUMP to Dump a Processor to Disk.....	135
Before You Begin.....	135
Dumping a Processor Using RCVDUMP (Excludes NonStop NS16000 Series and NS14000 Series Systems).....	135
Dumping a Processor Using RCVDUMP for NonStop NS16000 Series and NS14000 Series Systems.....	136
Backing Up a Processor Dump to Tape.....	137
Replacing Processor Memory.....	137
Replacing the Processor Board and Processor Entity.....	137
Submitting Information to Your Service Provider.....	138
Submitting Tapes of Processor Dumps.....	138
Submitting Tapes of Configuration and Operations Files.....	138
Additional Information Required by Your Service Provider	138
Related Reading for Processors.....	139
12 Disk Drives: Monitoring and Recovery.....	140
When to Use This Chapter.....	140
Overview of Disk Drives	140
Internal SCSI Disk Drives.....	141
M8xxx Fibre Channel Disk Drives.....	141
Enterprise Storage System (ESS) Disks.....	142
Serial Attached SCSI (SAS) Disks and Solid State Drives.....	142
Monitoring Disk Drives.....	143
Monitoring Disk Drives With OSM.....	143
Monitoring Drives With SCF	145
Examples.....	146
Monitoring the State of Disk Drives.....	148

Monitoring the Use of Space on a Disk Volume.....	149
Monitoring the Size of Database Files.....	149
Example.....	149
Monitoring Disk Configuration and Performance.....	150
Identifying Disk Drive Problems.....	150
SCSI Disk Drives Attached to NonStop NS-Series Systems.....	150
M8xxx Fibre Channel Disk Drives.....	151
Recovery Operations for Disk Drives	151
Recovery Operations for a Down Disk or Down Disk Path.....	152
Recovery Operations for a Nearly Full Database File.....	153
Related Reading for Disk Drives.....	154
13 Tape Drives: Monitoring and Recovery.....	155
When to Use This Chapter.....	155
Overview of Tape Drives.....	155
Monitoring Tape Drives.....	155
Monitoring Tape Drive Status With OSM.....	156
Monitoring Tape Drive Status With SCF.....	157
Example.....	158
Monitoring Tape Drive Status With MEDIACOM.....	158
Example.....	158
Monitoring the Status of Labeled-Tape Operations.....	159
Identifying Tape Drive Problems.....	159
Recovery Operations for Tape Drives	160
Recovery Operations Using the OSM Service Connection.....	160
Performing an OSM Action on a Tape Drive.....	160
Performing an OSM Action on a Multiple Tape Drives.....	160
Recovery Operations Using SCF.....	160
Related Reading for Tapes and Tape Drives.....	161
14 Printers and Terminals: Monitoring and Recovery.....	162
When to Use This Chapter.....	162
Overview of Printers and Terminals	162
Monitoring Printer and Collector Process Status.....	162
Monitoring Printer Status.....	162
Monitoring Collector Process Status.....	163
Recovery Operations for Printers and Terminals.....	163
Recovery Operations for a Full Collector Process.....	163
Related Reading for Printers.....	163
15 Applications: Monitoring and Recovery.....	165
When to Use This Chapter.....	165
Monitoring TMF.....	165
Monitoring the Status of TMF.....	165
Monitoring Data Volumes.....	166
TMF States.....	167
Monitoring the Status of Pathway.....	167
PATHMON States.....	168
Related Reading for Pathway.....	169
16 Power Failures: Preparation and Recovery.....	170
When to Use This Chapter.....	170
System Response to Power Failures	170
OSM Power Fail Support.....	170
NonStop Cabinets (Modular Cabinets)	171
NonStop S-Series I/O Enclosures.....	171
External Devices	171

ESS Cabinets.....	172
Networking CLIMs.....	172
Storage CLIMs.....	172
SAS Disk Drive Enclosures.....	172
Air Conditioning.....	172
Preparing for Power Failure.....	172
Set Ride-Through Time.....	173
Configure OSM Power Fail Support.....	173
Monitor Power Supplies.....	173
Monitor Batteries.....	173
Maintain Batteries.....	173
Power Failure Recovery.....	174
Procedure to Recover From a Power Failure.....	174
Setting System Time.....	174
Related Reading for Power Failures.....	174
17 Starting and Stopping the System	175
When to Use This Chapter.....	176
Powering On a System	176
Powering On the System From a Low-Power State.....	176
Powering On the System From a No Power State.....	177
Power Cycling a NonStop BladeSystem Processor.....	179
Power Cycling a NonStop NS2000 Series, NS2100, NS2200 Series, NS2300, or NS2400 Series Processor.....	181
Starting a System.....	182
Loading the System.....	183
Alerts.....	183
Normal System Load.....	183
System Load to a Specific Processor.....	183
System Load Disks.....	183
System Load Paths for a Normal System Load.....	184
Configuration File.....	185
Starting Other System Components.....	185
Performing a System Load	186
Performing a System Load From a Specific Processor.....	187
Reloading Processors.....	188
Reloading Processors Using the RELOAD Command.....	188
Reloading Processors Using OSM.....	189
Minimizing the Frequency of Planned Outages.....	189
Anticipating and Planning for Change.....	190
Stopping Applications, Devices, and Processes.....	190
Stopping the System.....	191
Alerts.....	191
Halting All Processors Using OSM.....	191
Powering Off a System.....	192
System Shutdown Using OSM.....	192
System Shutdown Using SCF.....	193
Emergency Power-Off Procedure.....	193
Troubleshooting and Recovery Operations	193
Fans Are Not Turning.....	193
System Does Not Appear to Be Powered On.....	193
Green LED Is Not Lit After POSTs Finish.....	194
Amber LED on a Component Remains Lit After the POST Finishes.....	194
Components Fail When Testing the Power	194
Recovering From a System Load Failure	194

Getting a Corrupt System Configuration File Analyzed.....	195
Recovering From a Reload Failure	195
Exiting the OSM Low-Level Link.....	196
Opening Startup Event Stream and Startup TACL Windows.....	196
Related Reading for Starting and Stopping a System.....	197
18 Creating Startup and Shutdown Files.....	199
Automating System Startup and Shutdown.....	200
NonStop Cluster Boot Application.....	200
Managed Configuration Services (MCS).....	200
Startup.....	201
Shutdown.....	201
For More Information	201
Processes That Represent the System Console.....	201
\$YMIOP.#CLCI.....	201
\$YMIOP.#CNSL.....	201
\$ZHOME.....	202
\$ZHOME Alternative.....	202
Example Command Files.....	202
CIIN File.....	203
Establishing a CIIN File.....	203
Modifying a CIIN File.....	204
If a CIIN File Is Not Specified or Enabled in OSM.....	204
Example CIIN Files.....	204
Writing Efficient Startup and Shutdown Command Files.....	206
Command File Syntax.....	206
Avoid Manual Intervention.....	206
Use Parallel Processing.....	206
Investigate Product-Specific Techniques.....	207
How Process Persistence Affects Configuration and Startup.....	207
Tips for Startup Files.....	207
Startup File Examples.....	208
System Startup File.....	208
Spooler Warm-Start File	209
TMF Warm-Start File.....	210
TCP/IP Stack Configuration and Startup File.....	210
CP6100 Lines Startup File.....	212
X.25 Lines Startup File.....	213
Expand-Over-IP Line Startup File.....	213
Expand Direct-Connect Line Startup File.....	214
Tips for Shutdown Files.....	214
Shutdown File Examples.....	214
System Shutdown File.....	214
CP6100 Lines Shutdown File.....	215
ATP6100 Lines Shutdown File.....	215
X.25 Lines Shutdown File.....	216
Printer Line Shutdown File.....	216
Expand-Over-IP Line Shutdown File.....	216
Direct-Connect Line Shutdown File.....	216
Spooler Shutdown File.....	217
TMF Shutdown File.....	217
19 Preventive Maintenance.....	218
When to Use This Chapter.....	218
Monitoring Physical Facilities.....	218
Checking Air Temperature and Humidity.....	218

Checking Physical Security.....	218
Maintaining Order and Cleanliness.....	218
Checking Fire-Protection Systems.....	219
Cleaning System Components.....	219
Cleaning an Enclosure.....	219
Cleaning and Maintaining Printers.....	219
Cleaning Tape Drives.....	219
Handling and Storing Cartridge Tapes.....	220
A Operational Differences Between Systems Running G-Series, H-Series, and J-Series RVUs.....	221
B Tools and Utilities for Operations.....	222
When to Use This Appendix.....	222
BACKCOPY.....	222
BACKUP.....	222
CLIMCMD.....	222
CLIMCONFIG.....	222
Disk Compression Program (DCOM).....	222
Disk Space Analysis Program (DSAP).....	222
EMSDIST.....	222
Event Management Service Analyzer (EMSA).....	222
File Utility Program (FUP).....	223
HP Systems Insight Manager (HP SIM).....	223
Insight Control Power Management.....	223
Insight Remote Support Advanced.....	223
Measure.....	223
MEDIACOM.....	223
NonStop Cluster Essentials.....	223
NonStop Cluster Performance Essentials.....	223
NonStop I/O Essentials.....	223
NonStop Maintenance LAN DHCP DNS Configuration Wizard.....	224
NonStop NET/MASTER.....	224
NonStop Software Essentials.....	224
NSKCOM and the Kernel-Managed Swap Facility (KMSF).....	224
Onboard Administrator (OA).....	224
OSM Product Suite.....	224
PATHCOM.....	224
PEEK.....	224
RESTORE.....	225
SPOOLCOM.....	225
Subsystem Control Facility (SCF)	225
HP Tandem Advanced Command Language (TACL).....	225
HP Tandem Failure Data System (TFDS).....	225
TFDSCOM.....	225
TMFCOM.....	226
Web ViewPoint	226
ViewSys.....	226
C Related Reading for Tools and Utilities.....	227
D Converting Numbers.....	232
When to Use This Appendix.....	232
Overview of Numbering Systems.....	232
Binary to Decimal.....	232
Example.....	232
Octal to Decimal.....	233

Example.....	233
Hexadecimal to Decimal.....	234
Example.....	235
Decimal to Binary.....	235
Example.....	235
Decimal to Octal.....	236
Example.....	236
Decimal to Hexadecimal.....	236
Example.....	237
Safety and Compliance.....	238
Index.....	243

About This Document

This guide describes how to perform routine system hardware operations for HP Integrity NonStop™ NS-series systems and HP Integrity NonStop BladeSystems. It covers H-series release version updates and J-series release version updates.

This guide is primarily geared toward commercial type NonStop BladeSystems and NonStop NS-series systems (see [“Differences Between NonStop NS-Series Systems”](#) (page 36) for high-level architectural and hardware differences between the various commercial models). While basic monitoring principles, such as [“Using the OSM Service Connection”](#) (page 59), apply to Telco as well as commercial systems, have your service provider refer to the carrier grade sections in the hardware installation manual for your system for hardware details and service procedures specific to Telco systems.

NOTE: NonStop BladeSystem refers to the hardware that makes up the NonStop BladeSystem. NonStop NS-series refers to the hardware that makes up the NonStop NS-series system.

The term *NonStop server* refers to NonStop NS-series servers, NonStop server blades in NonStop BladeSystems, and some NonStop NS-series systems.

H-series and J-series refer to the software that runs on NonStop BladeSystems and NonStop NS-series systems.

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

Supported Release Version Updates (RVUs)

This manual supports J06.03 and all subsequent J-series RVUs and H06.13 and all subsequent H-series RVUs until otherwise indicated in a replacement publication.

Intended Audience

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.

New and Changed Information in 529869-023 Edition

- Added information about the NonStop NS2400 series system and NS2300 system throughout the guide.
- Updated [“Powering Off a System”](#) (page 192).

New and Changed Information in 529869-022 Edition

- In Chapter 2 under [“NonStop BladeSystems NB50000c, NB50000c-cg, NB54000c, NB54000c-cg, NB56000c, and NB56000c-cg Modular Hardware Components”](#) (page 32), changed title and text to add the NonStop BladeSystems NB56000c and NB56000c-cg.
- Added references to and information about the NonStop BladeSystems NB56000c and NB56000c-cg as needed throughout the guide.

- Updated the chapter [“Overview of Monitoring and Recovery”](#) (page 51) to document other HP manageability tools and update information, as needed, for previously-documented tools.
- Added the chapter [“Monitoring the Performance of NonStop Subsystems”](#) (page 71) to document performance-monitoring tools.
- Updated the section [“Cluster I/O Modules \(CLIMs\)”](#) (page 87) to describe the OSM Service Connection System CLIM operations that generate CLIM diagnostic data.
- Updated [“Dumping a Processor to Disk”](#) (page 134) to include all supported NonStop systems and procedures.
- Updated [“Powering Off a System”](#) (page 192) to clarify the procedure.
- Added [“NonStop Cluster Boot Application”](#) (page 200) under [“Automating System Startup and Shutdown”](#) (page 200).
- Updated existing information and added new HP products, including TFDS and TFDSCOM, to [“Tools and Utilities for Operations”](#) (page 222).
- Added a reference to the *HP Tandem Failure Data System (TFDS) Manual* in [“Related Reading for Tools and Utilities”](#) (page 227).

New and Changed Information in 529869–021 Edition

- Added NonStop NS2100 system throughout the document where NonStop NS2200 series is discussed.
- Added Power Regulator in [“NonStop BladeSystem Options”](#) (page 33).
- Rearranged sections in [“Differences Between NonStop NS-Series Systems”](#) (page 36), added section [“NonStop NS2100 Systems”](#) (page 36), and changed [“NonStop NS2200 Series Systems”](#) (page 36) and [“NonStop NS2000 Series Systems”](#) (page 36).
- Changed [“Blade”](#) (page 39) to simplify it and add NonStop NS2100 systems.
- Added setting Power Regulator mode for NonStop BladeSystem NB54000c, and NB54000c-cg only in [“Insight Control Power Management”](#) (page 58).
- Changed section title of [“NonStop BladeSystem and NS2000 and NS2200 Series Processors”](#) to [“Processors in NonStop Systems Running J-Series RVUs”](#) (page 121) and added NonStop NS2100 processors.
- Added NonStop NS2100 blade element to [Table 23](#) (page 121).
- Added [“Managing and Setting NonStop BladeSystems NB54000c, NB54000c-cg, NB56000c, and NB56000c-cg Processor Power Usage with Insight Control Power Management”](#) (page 127).
- Changed title of reference to *Configuring NonStop NS2200 Series and NS2100 Systems and NonStop BladeSystems to Avoid System Startup Problems* in [“When to Use This Chapter”](#) (page 170).
- Changed title of [“Power Cycling a NonStop NS2000 or NS2100 Series Processor”](#) to [“Power Cycling a NonStop NS2000 Series, NS2100, NS2200 Series, NS2300, or NS2400 Series Processor”](#) (page 181).
- Changed title of reference to *Troubleshooting Tips for Powering On NonStop NS2100 Systems, NS2200 Series Systems and NonStop BladeSystems* in [“Recovering From a System Load Failure ”](#) (page 194).
- Added [“Insight Control Power Management”](#) (page 223).
- Added *HP SIM for NonStop Manageability* and Insight Control Power Management to [“Related Reading for Tools and Utilities”](#) (page 227).

New and Changed Information in 529869–020 Edition

- Changed note about OSM Notification Director and made other changes in [“System Consoles”](#) (page 28).
- Removed OSM Notification Director from lists and made other changes in [“Overview of OSM Applications”](#) (page 29) and [“Launching OSM Applications”](#) (page 30).
- Added core licensing to [“NonStop BladeSystem Options”](#) (page 33).
- Added link to [“NonStop Software Essentials”](#) (page 57) at the beginning of [“Overview of Monitoring and Recovery”](#) (page 51).
- Changed [“HP Insight Remote Support Advanced”](#) (page 59).
- Changed [“NonStop I/O Essentials”](#) (page 57) and removed list of configuration and control tasks in the CIP subsystem.
- Changed [“NonStop Software Essentials”](#) (page 57).
- Changed last item in [Table 12: “Related Reading for Monitoring Information”](#) (page 70).
- Changed [“Processors in NonStop Systems Running J-Series RVUs”](#) (page 121).
- Changed [Table 23](#) (page 121) and added paragraph after it about determining or increasing the number of BladeSystem cores enabled.
- For NonStop NB54000c and NB54000c-cg BladeSystems, added the Physical Cores attribute and Number of Enabled Cores attribute of the Logical Processor object. in [“Monitoring Processor Status Using the OSM Service Connection”](#) (page 125). Also told how to increase the number of cores enabled in a processor.
- Changed RVU requirements to support SSDs to H06.24 or later or J06.13 or later in [“Serial Attached SCSI \(SAS\) Disks and Solid State Drives”](#) (page 142) and added RVU requirement of H06.23 or later or J06.12 or later to support disk partitioning of SSDs or HDDs.
- Changed titles of references in [“Related Reading for Disk Drives”](#) (page 154).
- Changed [“NonStop Cluster Essentials”](#) (page 223).
- Changed [“CLIM Boot Service Configuration Wizard”](#) to [“NonStop Maintenance LAN DHCP DNS Configuration Wizard”](#) in Appendix B [“Tools and Utilities for Operations”](#) (page 222) and Appendix C [“Related Reading for Tools and Utilities”](#) (page 227).
- Changed [“NonStop Software Essentials”](#) (page 224).
- Added Down System CLIM Firmware Update Tool and made other changes in Appendix C [“Related Reading for Tools and Utilities”](#) (page 227).

New and Changed Information in 529869–019 Edition

- Under [“About This Document”](#) (page 13), changed the text in the second paragraph to genericize the manuals referred to, since information about NonStop BladeSystems, NS2000, and NS2200 carrier grade and commercial is combined in their hardware installation manuals.
- Under [“Related Information”](#) (page 23), changed text to refer to planning guides for CLIMs, rather than just for IP CLIMs and Storage CLIMs.
- In Chapter 2 under [“NonStop BladeSystems NB50000c, NB50000c-cg, NB54000c, NB54000c-cg, NB56000c, and NB56000c-cg Modular Hardware Components”](#) (page 32), changed title and text to add the NB50000c-cg and NB54000c-cg systems.
- In Chapter 2 under [“NonStop BladeSystems NB50000c, NB50000c-cg, NB54000c, NB54000c-cg, NB56000c, and NB56000c-cg Modular Hardware Components”](#) (page 32), added InfiniBand CLuster I/O Module (IB CLIM).
- In Chapter 2 under [“NonStop NS-Series Modular Hardware Components”](#) (page 34), added InfiniBand CLuster I/O Module (IB CLIM).

- In Chapter 2 under [“NonStop NS16000 Series Systems” \(page 37\)](#), indicated that NS16000 series servers running H06.20 or later RVUs support IP, Telco, and Storage CLIMs. NonStop NS16000 and NS16200 servers running H06.23 or later RVUs support IB CLIMs.
- In Chapter 2 under [“NonStop NS2000 Series Systems” \(page 36\)](#), added Telco to the list of CLIMs supported.
- In Chapter 2 under [“NonStop NS1000 and NS1200 Systems” \(page 38\)](#), added Telco to the list of CLIMs not supported.
- In Chapter 2, added the [“NonStop NS2200 Series Systems” \(page 36\)](#) section to describe the new NS2200 system.
- In Chapter 2 under [“Server and System” \(page 38\)](#), added references to the new NS2200 system.
- In Chapter 2 under [“Blade” \(page 39\)](#), added references to the new NS2200 system.
- In Chapter 2, made these changes to [“NonStop NS-Series Modular Hardware Components” \(page 34\)](#):
 - The Logical Synchronization Unit (LSU) applies to NonStop NS16000 series and NS14000 series systems only.
 - The Processor Switch, or p-switch applies to NonStop NS16000 series systems only.
 - The I/O Adapter Module (IOAM) Enclosure applies to NonStop NS16000 series, NS14000 series, and NS1000 systems only.
 - Added support for the NonStop NS2200 systems to the IP, Telco, and Storage CLIMs.
- In Chapter 3 under [“Summary of Monitored NonStop Hardware Resources” \(page 53\)](#), changed “IP CLIMs” to “CLIMs” to indicate all CLIMs.
- Effective with the H06.23/J06.12 RVU, solid state drives (SSDs) in SAS enclosures connected to G6 CLIMs are supported and are serviced and managed by OSM. The SMARTSSD Wear Gauge enables you to monitor the wear status of the drive.
Also, effective with the H06.23/J06.12 RVU, all SDDs and HDDs in SAS enclosures connected to G6 CLIMs can be partitioned. See the *SCF Reference Manual for the Storage Subsystem* for information about disk partitioning.
- In Chapter 3, added new section, [“Insight Control Power Management” \(page 58\)](#), describing this new HP SIM plug-in that provides the ability to visualize the layout of servers.
- In Chapter 6 under [“Local Area Networks \(LANs\) and Wide Area Networks \(WANs\)” \(page 85\)](#), made these changes:
 - Added statement: “An IB CLIM contains three customer-usable Ethernet interfaces and two customer-usable InfiniBand interfaces.”
 - Added note stating that IB CLIMs do not provide general purpose IB connectivity for NonStop systems and are only used as part of a Low Latency Solution.
 - Added note stating that all CLIMs and other I/O devices also connect to a dedicated service LAN (also called a maintenance LAN). For more information about the dedicated service LAN, ask your service provider to refer to the *NonStop Dedicated Service LAN Installation and Configuration Guide*
 - In the same section, added statement: “A Telco CLIM supports the Message Transfer Part Level 3 User Adaptation layer (M3UA) protocol.”

- In Chapter 7 under [“ServerNet and System I/O Connectivity” \(page 100\)](#):
 - Added IB CLIM and Telco CLIM to NonStop BladeSystem
 - Added IB CLIM and Telco CLIM to NonStop NS16000 series systems
 - Added Telco CLIM to NonStop NS2000 series system
 - Added references to the new NS2200 system
- In Chapter 7 under [“When to Use This Chapter” \(page 100\)](#), added references to the new NS2200 system.
- In Chapter 7 under [“Monitoring the ServerNet Fabrics Using OSM” \(page 102\)](#), added references to the new NS2200 system and modified text concerning ServerNet connectivity.
- In Chapter 9, referred to the *Cluster I/O Protocols (CIP) Configuration and Management Manual* for information about shutting down a CLIM. Added this reference to [“Networking CLIMs” \(page 172\)](#) in Chapter 16, as well.
- In Chapter 10 under [“Processors in NonStop Systems Running J-Series RVUs” \(page 121\)](#), added references to the new NS2200 system.
Indicated that the NonStop BladeSystem server blade and NonStop NS2000 series server blade utilize a two-core microprocessor.
- In Chapter 10 under [“Monitoring Processor Status Using the OSM Service Connection” \(page 125\)](#), added references to the new NS2200 system.
- In Chapter 10 under [“Recovery Operations for Processors” \(page 129\)](#), added references to the new NS2200 system and changed introductory text for systems supported.
- In Chapter 11 under [“Serial Attached SCSI \(SAS\) Disks and Solid State Drives” \(page 142\)](#), added paragraph:
Effective with the H06.23/J06.12 RVU, solid state drives (SSDs) in SAS enclosures connected to G6 CLIMs are supported. SSDs, as well as HDDs in SAS enclosures connected to G6 CLIMs, can be partitioned. OSM provides support for SSDs and displays partitioning information for HDDs and SSDs. See the *SCF Reference Manual for the Storage Subsystem* for more information.
- In Chapter 11 under [“Enterprise Storage System \(ESS\) Disks” \(page 142\)](#), changed the term “StorageWorks” to storage due to rebranding.
- In Chapter 11 under [“Common Recovery Operations for Disk Drives ” \(page 151\)](#), indicated that you should refer to the installation manual for your system for instructions on how to create an alternate system disk.
- In Chapter 15 under [“Networking CLIMs” \(page 172\)](#), added Telco CLIMs and IB CLIMs to this section describing how power failures and recoveries affect these CLIMs.
- In Chapter 15 under [“Storage CLIMs” \(page 172\)](#), stated that a Storage CLIM with the Write Cache Enable option selected and a local UPS on that rack can preserve power long enough for the cached data to be flushed to disk. The WRITECACHE ENABLED option of the ADD DISK command is described in the *SCF Reference Manual for the Storage Subsystem*.
- In Chapter 15 under [“SAS Disk Drive Enclosures” \(page 172\)](#), added paragraph:
All disk drives in SAS enclosures connected to G6 CLIMs, including any partitions they might have, are supported by OSM. OSM will alarm and do a dialout when a SAS solid state drive (SSD) approaches or has reached its wear-out threshold.

- In Chapter 16, added references to the new NS2200 system under:
 - [“Powering On the System From a Low-Power State”](#) (page 176)
 - [“Powering On the System From a No Power State”](#) (page 177)
 - [“Starting a System”](#) (page 182)
 - [“System Load Disks”](#) (page 183)
 - [“Powering Off a System”](#) (page 192)
- Changed and added steps for both NS2000 and NS2200 systems under [“Power Cycling a NonStop NS2000 Series, NS2100, NS2200 Series, NS2300, or NS2400 Series Processor”](#) (page 181).
- In Chapter 16 under [“Starting a System”](#) (page 182), changed text to indicate that only the NonStop BladeSystems require a Halted State Services (HSS) firmware file on the dedicated service LAN.
- In Chapter 16 under [“System Load Disks”](#) (page 183) and [“Related Reading for Starting and Stopping a System”](#) (page 197), indicated that you should refer to the installation manual for your system for instructions on how to create an alternate system disk.

New and Changed Information in 529869–018 Edition

This edition adds a table of contents to the beginning of the manual.

New and Changed Information in 529869–017 Edition

This edition includes support for the new NonStop NB54000c BladeSystem, which has a quad-core microprocessor.

Change bars mark the following information, which has been added to or changed for this version of the manual:

- Added NB54000c BladeSystems to [“Determining Your System Configuration”](#) (page 32) and [“NonStop BladeSystems NB50000c, NB50000c-cg, NB54000c, NB54000c-cg, NB56000c, and NB56000c-cg Modular Hardware Components”](#) (page 32) in Chapter 2.
- Added the Telco CLuster I/O Module (Telco CLIM) and the HP CLIM Cable Mangement Ethernet patch panel to [“NonStop BladeSystems NB50000c, NB50000c-cg, NB54000c, NB54000c-cg, NB56000c, and NB56000c-cg Modular Hardware Components”](#) (page 32) in Chapter 2.
- Changed text for the Serial attached SCSI (SAS) disk drive enclosure and added a reference note under [“NonStop BladeSystems NB50000c, NB50000c-cg, NB54000c, NB54000c-cg, NB56000c, and NB56000c-cg Modular Hardware Components”](#) (page 32) in Chapter 2.
- Changed references for the IP CLuster I/O Module (IP CLIM) and the Storage CLuster I/O Module (Storage CLIM) under [“NonStop NS-Series Modular Hardware Components”](#) (page 34) in Chapter 2.
- Added the Telco CLuster I/O Module (Telco CLIM) and the HP CLIM Cable Mangement Ethernet patch panel to [“NonStop NS-Series Modular Hardware Components”](#) (page 34) in Chapter 2.
- Changed text for the Serial attached SCSI (SAS) disk drive enclosure and added a reference note under [“NonStop NS-Series Modular Hardware Components”](#) (page 34) in Chapter 2.
- Added quad-core microprocessor to [“Processors in NonStop Systems Running J-Series RVUs”](#) (page 121) in Chapter 10.
- Changed text for SAS disk drives under [“Serial Attached SCSI \(SAS\) Disks and Solid State Drives”](#) (page 142) in Chapter 11.

- Changed reference under “Overview of Tape Drives” (page 155), “Related Reading for Disk Drives” (page 154) and “Related Reading for Tapes and Tape Drives” (page 161) in Chapter 12.
- Added note pointing to reference on configuration settings for NonStop BladeSystems that shipped before J06.10 under “When to Use This Chapter” (page 170) in Chapter 15.
- Added new text to “Starting a System” (page 182) in Chapter 16.
- Added a new reference to “Recovering From a System Load Failure ” (page 194) in Chapter 16.
- Corrected text in “NonStop NS14000 Series Systems” (page 37) and “4-Port ServerNet Extender (4PSE)” (page 108) to state that FCSAs and G4SAs can only be installed in IOAM enclosure slots 3 through 5.

New and Changed Information in 529869–016 Edition

The major changes in this edition are support for Storage CLIMs on NonStop NS16000 series systems and support for volume level encryption on NonStop BladeSystems and NonStop NS16000 series systems. Also, HP Insight Remote Support Advanced is added as go-forward replacement for OSM Notification Director. and the HP NonStop Software Essentials plug-in for HP SIM was added.

Change bars mark the following information, which has been added to or changed for this version of the manual:

- Changed document title to *NonStop Operations Guide for H-Series and J-Series RVUs*.
- Changed URLs for NTL (external), NTL (internal), and WOW in “Related Information” (page 23). Also changed titles to *Tips for Locating NonStop Documents in the BSC* and *Tips for Locating NonStop Documents in SAW* and changed the URLs.
- Changed URLs for NTL (external), NTL (internal), and WOW in “Support and Service Collection” (page 31). Also changed titles to *Tips for Locating NonStop Documents in the BSC* and *Tips for Locating NonStop Documents in SAW* and changed the URLs.
- Added note to “System Consoles” (page 28) that OSM Notification Director functionality has been replaced by HP Insight Remote Support Advanced.
- Changed description of OSM Notification Director in “Overview of OSM Applications” (page 29).
- Added note to “Launching OSM Applications” (page 30) about OSM Notification Director not visible when using Remote Desktop Connection to access a remote NonStop system console.
- Added volume level encryption to “NonStop BladeSystem Options” (page 33) and “NonStop NS16000 Series Systems” (page 37).
- Added IP CLIM, Storage CLIM, and MSA70 SAS disk drive enclosure to “NonStop NS-Series Modular Hardware Components” (page 34).
- Changed “Monitoring Problem Incident Reports” (page 59) to add Insight Remote Support Advanced.
- Added “Manage” to title of “Using HP SIM” (page 55), added NonStop Software Essentials to list there of HP SIM plug-ins, and added pointer to general HP SIM manuals.
- Added “HP Insight Remote Support Advanced” (page 59).
- Changed description of single-click boot application in “NonStop Cluster Essentials” (page 58).
- Changed first bullet list item in “NonStop Cluster Performance Essentials” (page 58).
- Added “Configure and control storage encryption” and “Configure and control Telco INS subsystem” to “NonStop I/O Essentials” (page 57).

- Added “NonStop Software Essentials” (page 57).
- Removed “HP Storage Essentials”.
- Added OSM Service Connection online help to Table 12 (page 70).
- Changed “Monitoring the CLIMs with HP SIM” (page 112).
- Added references for monitoring and managing CLIMs to “Related Reading for CLIMs” (page 119).
- Changed Page 173 for NonStop BladeSystems to configure OSM power fail support for only Enclosure 100.
- Added “NonStop Software Essentials” (page 224).

New and Changed Information in 529869–015 Edition

The major change in this edition is support for NonStop NS16000 series systems as nodes in the BladeCluster Solution topology. Also, BladeClusters can co-exist with 6770 ServerNet clusters and 6780 ServerNet clusters. For hardware and software requirements, see the *BladeCluster Solution Manual*.

- The following have been added to the note in “When to Use This Chapter” (page 100) in Chapter 8 (page 100):
 - BladeCluster Solution availability for NonStop NS16000 series systems
 - Co-existence of BladeClusters with 6770 ServerNet clusters or 6780 ServerNet clusters
- The *BladeCluster Solution Manual* and BladeCluster ServerNet connectivity for NonStop NS16000 series systems have been added to Table 18 (page 100).
- The Safety and Compliance statement has been updated at the end of this manual.

Document Organization

This document is organized as follows:

- “Introduction to NonStop Operations” (page 26)
- “Determining Your System Configuration” (page 32)
- “Overview of Monitoring and Recovery” (page 51)
- “Monitoring the Performance of NonStop Subsystems” (page 71)
- “Monitoring EMS Event Messages” (page 77)
- “Processes: Monitoring and Recovery” (page 79)
- “Communications Subsystems: Monitoring and Recovery” (page 85)
- “ServerNet Resources: Monitoring and Recovery” (page 100)
- “I/O Adapters and Modules: Monitoring and Recovery” (page 106)
- “Cluster I/O Modules (CLIMs): Monitoring and Recovery” (page 112)
- “Processors and Components: Monitoring and Recovery” (page 120)
- “Disk Drives: Monitoring and Recovery” (page 140)
- “Tape Drives: Monitoring and Recovery” (page 155)
- “Printers and Terminals: Monitoring and Recovery” (page 162)
- “Applications: Monitoring and Recovery” (page 165)
- “Power Failures: Preparation and Recovery” (page 170)

- “Starting and Stopping the System ” (page 175)
- “Creating Startup and Shutdown Files” (page 199)
- “Preventive Maintenance” (page 218)
- “Operational Differences Between Systems Running G-Series, H-Series, and J-Series RVUs” (page 221)
- “Tools and Utilities for Operations” (page 222)
- “Related Reading for Tools and Utilities” (page 227)
- “Converting Numbers” (page 232)
- “Safety and Compliance” (page 238)

Notation Conventions

General Syntax Notation

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

UPPERCASE LETTERS

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

MAXATTACH

Italic Letters

Italic letters, regardless of font, indicate variable items that you supply. Items not enclosed in brackets are required. For example:

file-name

Computer Type

Computer type letters indicate:

- C and Open System Services (OSS) keywords, commands, and reserved words. Type these items exactly as shown. Items not enclosed in brackets are required. For example:

Use the `cextdecs.h` header file.

- Text displayed by the computer. For example:

Last Logon: 14 May 2006, 08:02:23

- A listing of computer code. For example

```
if (listen(sock, 1) < 0)
{
  perror("Listen Error");
  exit(-1);
}
```

Bold Text

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

ENTER RUN CODE

?123

CODE RECEIVED: 123.00

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

[] 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 can 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 can 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 typed 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 type 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 this example, no spaces are 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 ]...
```

Related Information

Operations planning and operations management practices appear in these manuals:

- The planning guide for your NonStop system
- *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, which is now on www.hp.com/go/nonstop-docs.

For comprehensive information about performing operations tasks for a NonStop BladeSystem or for a NonStop NS-series system, 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 systems. 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.

For information on how install, configure, and use OSM server and client components:

- *OSM Configuration Guide*
- *NonStop System Console Installer Guide*
- *OSM Service Connection User's Guide* (online help or NTL version)

Systems that are connected in ServerNet clusters (such as a 6770 ServerNet clusters, 6780 ServerNet clusters, or BladeClusters) require special installation, configuration, and operating procedures. For such information, depending on the ServerNet topology employed, see the appropriate manuals listed below.

For the BladeCluster Solution, see the *BladeCluster Solution Manual*.

In the 6780 ServerNet cluster environment, which uses the layered topology, see:

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

- *ServerNet Cluster Supplement for NonStop NS-Series Servers* or the *ServerNet Cluster Supplement for NonStop BladeSystems*
- The *Add Node to ServerNet Cluster* guided procedure online help

For the earlier ServerNet cluster star topologies, using 6770 switches, see:

- *ServerNet Cluster Manual*
- *ServerNet Cluster Supplement for NonStop NS-Series Servers* or the *ServerNet Cluster Supplement for NonStop BladeSystems*
- The *Add Node to ServerNet Cluster* guided procedure online help

For information about CLIMs, refer to the planning guide for your NonStop system.

For information about the system console, refer to the *NonStop System Console Installer Guide*.

For information about NonStop Cluster Essentials, refer to the *NonStop Cluster Essentials Installation and Quick Start Guide*.

For information about NonStop I/O Essentials, refer to the *NonStop I/O Essentials Installation and Quick Start Guide*.

For other documentation related to operations tasks, see “[Related Reading for Tools and Utilities](#)” (page 227).

Support and Service Collection

The NTL Support and Service collection provides procedures, part numbers, troubleshooting tips, and tools for servicing NonStop systems.

NOTE: NTL is now on www.hp.com/go/nonstop-docs (external view) and www.hp.com/go/nonstop-support-docs (internal view). Internal users should consult *Tips for Locating NonStop Documents in SAW* at http://h41302.www4.hp.com/km/saw/view.do?docId=emr_na-c02163441.

Authorized service providers can also order the NTL Support and Service Library CD:

- HP employees: Subscribe at World on a Workbench (WOW). Subscribers automatically receive CD updates. Access the WOW order form at http://teams1.sharepoint.hp.com/teams/WOW_Onsite/default.aspx.
- Channel Partners and Authorized Service Providers: contact pubs.comments@hp.com.

Publishing History

Part Number	Product Version	Publication Date
529869-022	N.A.	August 2013
529869-021	N.A.	August 2012
529869-020	N.A.	February 2012
529869-019	N.A.	August 2011
529869-018	N.A.	February 2011
529869-017	N.A.	February 2011

HP Encourages Your Comments

HP encourages your comments concerning this document. We are committed to providing documentation that meets your needs. Send any errors found, suggestions for improvement, or compliments to:

docsfeedback@hp.com

Include the document title, part number, and any comment, error found, or suggestion for improvement you have concerning this document.

1 Introduction to NonStop Operations

- “When to Use This Chapter” (page 26)
- “Understanding the Operational Environment” (page 26)
- “What Are the Operator Tasks?” (page 26)
- “Logging On to a NonStop System” (page 28)
 - “System Consoles” (page 28)
 - “Opening a TAACL Window” (page 29)
 - “Overview of OSM Applications” (page 29)
 - “Launching OSM Applications” (page 30)
- “Service Procedures” (page 31)
 - “Support and Service Collection” (page 31)

When to Use This Chapter

This chapter identifies operator tasks for NonStop system hardware and software components. It provides an introduction to the other chapters and sections in this guide that describe these operator tasks in more detail.

Understanding the Operational Environment

To understand the operational environment:

- If you are already familiar with other NonStop systems, see [Appendix A \(page 221\)](#).
- For a brief introduction to the system organization and the location of system components in an NonStop system, see [Chapter 2 \(page 32\)](#).
- For information about various software tools and utilities you can use to perform system operations on an NonStop system, see [Appendix B \(page 222\)](#).

What Are the Operator Tasks?

The system operator tasks described in this guide include:

- “Starting the NonStop System and Loading the NonStop Operating System” (page 27)
- “Updating Firmware” (page 27)
- “Configuring and Controlling System Components” (page 27)
- “Monitoring System Hardware and Performance Subsystems and Performing Recovery Operations ” (page 27)
- “Preparing for and Recovering from Power Failures” (page 28)
- “Stopping and Powering Off the System” (page 28)
- “Powering On and Starting the System” (page 28)
- “Performing Preventive Maintenance” (page 28)
- “Operating Disk Drives and Tape Drives” (page 28)
- “Responding to Spooler Problems” (page 28)

Starting the NonStop System and Loading the NonStop Operating System

For information about starting a NonStop system, see “Starting a System” (page 182).

For information about loading the NonStop operating system, see “Loading the System” (page 183) and the *Software Installation and Upgrade Guide* for your RVU.

Updating Firmware

For information about installing compatible versions of firmware, see the *Software Installation and Upgrade Guide* for your RVU and the *NonStop Firmware Matrices*.

Configuring and Controlling System Components

You can configure and control these NonStop components:

- For CLIMs, see the *Cluster I/O Protocols (CIP) Configuration and Management Guide*.
- For disks, see the *SCF Manual for the Storage Subsystem* or the management guide for a particular disk product.
- For tapes, see the *SCF Manual for the Storage Subsystem* or the management guide for a particular tape product.
- For the OSS environment, see *Open System Services Management and Operations Guide* for J-series or H-series RVUs.
- For Expand networks, see the *Expand Configuration and Management Manual*.
- For Backup jobs, see the *Backup and Restore 2 Manual*.

NOTE: TMF, Pathway, SQL/MX, and SQL/MP must also be configured and controlled, but these are tasks are not performed by system operators and are not documented in this guide.

Monitoring System Hardware and Performance Subsystems 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 NonStop systems can be performed online. Monitoring the status of all system components and performing recovery operations are described in these chapters:

- “Overview of Monitoring and Recovery” (page 51)
- “Monitoring the Performance of NonStop Subsystems” (page 71)
- “Monitoring EMS Event Messages” (page 77)
- “Processes: Monitoring and Recovery” (page 79)
- “Communications Subsystems: Monitoring and Recovery” (page 85)
- “ServerNet Resources: Monitoring and Recovery” (page 100)
- “I/O Adapters and Modules: Monitoring and Recovery” (page 106)
- “Processors and Components: Monitoring and Recovery” (page 120)
- “Disk Drives: Monitoring and Recovery” (page 140)
- “Tape Drives: Monitoring and Recovery” (page 155)
- “Printers and Terminals: Monitoring and Recovery” (page 162)
- “Applications: Monitoring and Recovery” (page 165)

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 planning guide for your NonStop system.

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 [Chapter 16 \(page 170\)](#).

Stopping and Powering Off the System

HP recommends a specific set of procedures for stopping a NonStop system and powering off its components, as described in [Chapter 17 \(page 175\)](#).

Powering On and Starting the System

HP recommends a specific set of procedures for powering on and starting a NonStop system or its components, as described in [Chapter 17 \(page 175\)](#).

Creating Startup and Shutdown Files

HP recommends a specific set of procedures for creating startup and shutdown files on a NonStop system or its components, as described in [Chapter 18 \(page 199\)](#).

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 [Chapter 19 \(page 218\)](#).

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*.

Logging On to a NonStop System

Many operations and troubleshooting tasks are performed by logging on to your NonStop system 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” \(page 29\)](#).

System Consoles

A system console is a personal computer approved by HP to run maintenance and diagnostic software for NonStop systems. 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 DVD.

System consoles communicate with NonStop systems 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.

NOTE: For information about configuring and enabling a remote desktop for dial-in services for use by a service provider or authorized employee, see the *NonStop System Console Installer Guide*.

The OSM Low-Level Link and OSM Console Tools components 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 system, and connectivity is established from the console through Internet Explorer browser sessions. For more information, see [“Launching OSM Applications” \(page 30\)](#).

NOTE: The OSM Notification Director functionality has been replaced by HP Insight Remote Support Advanced, which can be installed on the NonStop system consoles or on HP Central Management Servers (CMS) for [“Monitoring Problem Incident Reports” \(page 59\)](#). The OSM Notification Director may still reside on existing system consoles, but it is no longer installed on new consoles and it is not delivered on the NonStop system console installer DVD.

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 MR-Win6530

comForte MR-Win6530 is the terminal emulator shipped with the console software. It is launched by the OSM Low-Level Link version T0633AAT and later for Startup TACL and Event Stream windows. To launch a TACL session directly from MR-Win6530, see the MR-Win6530 User's Guide (shipped with the console-based product).

Opening a TACL Window From the OSM Low-Level Link

If the system has not started any TCP/IP connections, you can 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” \(page 196\)](#).

NOTE: OSM Low-Level Link version T0633AAT and later launches comForte MR-Win6530 for TACL sessions.

For more details on the functions of the TACL command interpreter, see [Appendix B \(page 222\)](#).

Overview of OSM Applications

The HP NonStop Open System Management (OSM) product suite, used to perform a variety of functions, includes:

- The OSM Service Connection is the primary interface for managing and servicing NonStop systems and components, as well as the ServerNet Cluster components (such as in 6770 ServerNet cluster, 6780 ServerNet cluster, or BladeCluster). See [“Using the OSM Service Connection” \(page 59\)](#) for an overview of how the OSM Service Connection is used to monitor your system components.
- The OSM Low-Level Link provides down-system support, such as [“Starting a System” \(page 182\)](#) and [“Recovery Operations for Processors” \(page 129\)](#), but also allows you to configure CLIMs, IOAM, VIO, and p-switch, and ServerNet switch modules, update CLIM software and firmware, and create or update HSS files necessary for the processors in an all J-series (except the NS2000) to boot.

- The OSM Event Viewer is used for “Monitoring EMS Event Messages” (page 77).
- OSM Console Tools:
 - NonStop Maintenance LAN DHCP DNS Configuration Wizard — used to configure DHCP, DNS, and BOOTP servers required for certain NonStop Systems.
 - Down System CLIM Firmware Update Tool — used for updating firmware/BIOS for CLIM components during planned system down time.
 - OSM Certificate Tool — used for HP Integrity NonStop BladeSystems to facilitate communication between OSM and the Onboard Administrator (OA) for those systems.
 - OSM System Inventory Tool — used to create an inventory file listing the hardware, firmware, or Numeric Sensor data for one or more NonStop systems running OSM.
 - Apache OpenOffice — a third-party product, installed as part of the OSM Console Tools, provides a method to view spreadsheets, documents, presentations of the system console; for example, files created by OSM in .csv (or comma delimited) format, such as hardware inventories created by the OSM System Inventory Tool and Access Control List reports or saved Multi-Resource Views created in the OSM Service Connection.
 - Terminal Emulator File Converter — used to convert OSM Service Connection-related OutsideView session files to MR:Win6530 format.

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>All Programs>HP OSM**. Then select the name of the application to launch:

- OSM Service Connection*
- OSM Low-Level Link Application
- OSM Event Viewer*
OSM System Inventory Tool
- OSM Certificate Tool
- Down System CLIM Firmware Update Tool
- NonStop Maintenance LAN DHCP DNS Configuration Wizard
- Terminal Emulator File Converter

NOTE: Apache OpenOffice is launched from **Start>All Programs>OpenOffice**

NOTE: If the OSM Notification Director is not visible when accessing a remote NonStop system console (NSC) using Microsoft's Remote Desktop Connection — and the OSM Notification Director is installed on that console and still being used for remote support instead of its replacement product, HP Insight Remote Support Advanced — you must start the session using a different command depending on the client version of Remote Desktop Connection. Refer to *OSM Notification Director Not Visible When Accessing Remote NonStop System Console* in Trouble Shooting Tips in the Support and Service Collection on NTL.

*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 Configuration Guide*
- *OSM Service Connection User's Guide* (NTL version of the online help)
- Online help within the OSM Service Connection, Low-Level Link, Event Viewer, and individual OSM Console Tools 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, see "[Support and Service Collection](#)" (page 31).

Support and Service Collection

The NTL Support and Service collection provides procedures, part numbers, troubleshooting tips, and tools for servicing NonStop systems.

NOTE: NTL is now on www.hp.com/go/nonstop-docs (external view) and www.hp.com/go/nonstop-support-docs (internal view). For more information, see *Tips for Locating NonStop Documents in the BSC* at <http://bizsupport1.austin.hp.com/bc/docs/support/SupportManual/c02141785/c02141785.pdf>. Internal users should also consult *Tips for Locating NonStop Documents in SAW* at http://h41302.www4.hp.com/km/saw/view.do?docId=emr_na-c02163441.

Authorized service providers can also order the NTL Support and Service Library CD:

- 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>.
- Channel Partners and Authorized Service Providers: contact pubs.comments@hp.com.

2 Determining Your System Configuration

- “When to Use This Chapter” (page 32)
- “NonStop BladeSystems” (page 32)
 - “NonStop BladeSystems NB50000c, NB50000c-cg, NB54000c, NB54000c-cg, NB56000c, and NB56000c-cg Modular Hardware Components” (page 32)
 - “NonStop BladeSystem Options” (page 33)
- “NonStop NS-Series Systems” (page 34)
 - “NonStop NS-Series Modular Hardware Components” (page 34)
 - “Differences Between NonStop NS-Series Systems” (page 36)
- “Terms Used to Describe System Hardware Components” (page 38)
- “Recording Your System Configuration” (page 39)
- “Using SCF to Determine Your System Configuration” (page 39)
 - “SCF System Naming Conventions” (page 40)
 - “SCF Configuration Files” (page 40)
 - “Using SCF to Display Subsystem Configuration Information” (page 41)
 - “Displaying SCF Configuration Information for Subsystems ” (page 43)
 - “Additional Subsystems Controlled by SCF” (page 47)
 - “Displaying Configuration Information—SCF Examples” (page 48)
- “CIP Subsystem Configuration” (page 50)

When to Use This Chapter

This chapter describes the system enclosures, the system organization, numbering and labeling, and how to identify components in a NonStop BladeSystem or a NonStop NS-series system. For detailed information about system hardware organization, refer to the appropriate planning guide for your NonStop system.

NonStop BladeSystems

NonStop BladeSystems NB50000c, NB50000c-cg, NB54000c, NB54000c-cg, NB56000c, and NB56000c-cg Modular Hardware Components

Hardware for NonStop BladeSystem NB50000c, NB50000c-cg, NB54000c, NB54000c-cg, NB56000c, and NB56000c-cg is implemented in modules or enclosures that are installed in modular cabinets. The systems include these hardware components:

- HP NonStop BladeSystem c7000 enclosure, which contains these components:
 - Server blades (up to 8)
 - Six power supplies for the NB50000c, NB54000c, or NB56000c; six DC power supplies for the NB50000c-cg, NB54000c-cg, or NB56000c-cg
 - Ten active cool fans

- Interconnect Ethernet switch for the NB50000c, NB54000c, or NB56000c; Interconnect Ethernet switch CG for the NB50000c-cg, NB54000c-cg, or NB56000c-cg
- Onboard Administrator (OA) module
- 1-phase or 3-phase AC input power module for the the NB50000c, NB54000c, or NB56000c; CD input module for the NB50000c-cg, NB54000c-cg, or NB56000c-cg
- 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)
- 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)
- IP CLuster I/O Module (IP CLIM)
- Telco CLuster I/O Module (Telco CLIM)
- HP CLIM Cable Mangement Ethernet patch panel
- InfiniBand CLuster I/O Module (IB CLIM)
- Storage CLuster I/O Module (Storage CLIM)
- Serial attached SCSI (SAS) disk drive enclosure (requires compatible Storage CLIM in system)

NOTE: For Storage CLIMs and SAS disk enclosures supported by your NonStop system, see the planning guide for your NonStop system.

For more information about these and other NonStop BladeSystem hardware components, see the *NonStop BladeSystem Planning Guide*.

NonStop BladeSystem Options

NonStop BladeSystems offer a variety of configuration options to suit different customer needs. For more information, see the *NonStop BladeSystem Planning Guide*.

Volume level encryption is supported on NonStop BladeSystems that are licensed to support it. If your system is licensed for volume level encryption, you can encrypt disk and tape volumes that are connected to the system with a Storage CLIM. See the *NonStop CLuster I/O Module (CLIM) Installation and Configuration Guide* for details.

Core licensing is supported on NonStop BladeSystems NB54000c, NB54000c-cg, NB56000c, and NB56000c-cg. The core license file can be updated to increase the number of microprocessor cores (IPUs) enabled in a processor. See the *NonStop BladeSystem Planning Guide* for details.

HP Power Regulator can provide centralized control of server power consumption on NonStop BladeSystems NB54000c, NB54000c-cg, NB56000c, and NB56000c-cg. To set up Power Regulator, first use the OSM Service Connection to enable Power Regulator, then immediately use Insight Control Power Management to set the Power Regulator mode. See *HP SIM for NonStop Manageability* for these procedures, Power Regulator modes, and minimum requirements.

Differences Between NonStop BladeSystems

NonStop BladeSystems NB50000c and NB50000c-cg

NonStop BladeSystems NB50000c and NB50000c-cg differ from the other BladeSystems in these respects:

- Run on J06.04 and later J-series RVUs
- Use from 8 GB to 48 GB main memory per logical processor

For more information, see the *NonStop BladeSystems Planning Guide* and *NonStop BladeSystems Hardware Installation Manual*.

NonStop BladeSystems NB54000c and NB55000c-cg

NonStop BladeSystems NB54000c and NB54000c-cg differ from the other BladeSystems in these respects:

- Run on J06.11 and later J-series RVUs
- Use from 16 GB to 64 GB main memory per logical processor

For more information, see the *NonStop BladeSystems Planning Guide* and *NonStop BladeSystems Hardware Installation Manual*.

NonStop BladeSystems NB56000c and NB56000c-cg

NonStop BladeSystems NB56000c and NB56000c-cg differ from the other BladeSystems in these respects:

- Run on J06.16 and later J-series RVUs
- Use from 16 GB to 96 GB main memory per logical processor
- Use BLC860c i4 Server Blades
- The NB56000c-cg supports the Gen8 CLIM CG
- New BladeSystems NB56000c and NB56000c-cg ship with an upgraded version of the C7000 enclosure. Field upgrades of an NB50000c or NB54000c to an NB56000c and NB50000c-cg or NB54000c-cg to an NB56000c-cg continue to use the earlier versions of the c7000 enclosure.

For more information, see the *NonStop BladeSystems Planning Guide* and *NonStop BladeSystems Hardware Installation Manual*.

NonStop NS-Series Systems

NonStop NS-Series Modular Hardware Components

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

- Modular Cabinet with Power Distribution Unit (PDU)
- NonStop Blade Complex
- NonStop Blade Element
- Logical Synchronization Unit (LSU) — applies to NonStop NS16000 series and NS14000 series systems only
- Processor Switch, or p-switch — applies to NonStop NS16000 series systems only

- I/O Adapter Module (IOAM) Enclosure — applies to NonStop NS16000 series, NS14000, and NS1000 systems only
IOAMs include these subcomponent I/O Adapters:
 - Fibre Channel ServerNet adapter (FCSA)
 - Gigabit Ethernet 4-port ServerNet adapter (G4SA)
 - 4-Port ServerNet Extenders (4PSEs) (NonStop NS14000 and NS1000 systems only)
- VIO Enclosure (displayed by OSM as a VIO Module object) — For more information, see [“NonStop NS14000 Series Systems” \(page 37\)](#), [“NonStop NS2000 Series Systems” \(page 36\)](#), [“NonStop NS1000 and NS1200 Systems” \(page 38\)](#), 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)
- IP CLuster I/O Module (IP CLIM) only in:
 - NonStop NS16000 series systems (refer to the *NonStop NS16000 Series Planning Guide* for CLIM models and RVU requirements)
 - NonStop NS2000 series, NonStop NS2100, NonStop NS2200 series, NonStop NS2300 systems, and NonStop NS2400 series systems
- Telco CLuster I/O Module (Telco CLIM) only in:
 - NonStop NS16000 series systems (refer to the *NonStop NS16000 Series Planning Guide* for CLIM models and RVU requirements)
 - NonStop NS2000 series systems, NonStop NS2100 systems, NonStop NS2200 series, NonStop NS2300 systems, and NonStop NS2400 series systems
- HP CLIM Cable Management Ethernet patch panel
- InfiniBand CLuster I/O Module (IB CLIM) only in:
 - NonStop NS16000 systems (refer to the *NonStop NS16000 Series Planning Guide* for CLIM models and RVU requirements)
- Storage CLuster I/O Module (Storage CLIM) only in:
 - NonStop NS16000 series systems (refer to the *NonStop NS16000 Series Planning Guide* for CLIM models and RVU requirements)
 - NonStop NS2000 series systems, NonStop NS2100 systems, NonStop NS2200 series, NonStop NS2300 systems, and NonStop NS2400 series systems
- Serial attached SCSI (SAS) disk drive enclosure (requires compatible Storage CLIM in system)

NOTE: For Storage CLIMs and SAS disk enclosures supported by your NonStop system, see the planning guide for your NonStop system.

NonStop NS-Series System Options

NonStop NS-series systems offer a variety of architecture and configuration options to suit different customer needs. For more information, see the appropriate planning guide for your NonStop system.

For information about supported CLIMs or CLIM-attached storage (SAS disk drive enclosures), see the planning guide for your NonStop system.

Differences Between NonStop NS-Series Systems

NonStop NS2400 Series Systems

NonStop NS2400 series systems are first released with J06.17. The NonStop NS2400 series systems (NS2400/T/ST) differ from the NS2200 series systems as follows: .

- NonStop NS2400 series systems use a different blade element .
- NonStop NS2400 series systems has different memory options.

For more information about NonStop NS2400 series systems, refer to the *NonStop NS2400 Series Planning Guide*

NonStop NS2300 Systems

The NonStop NS2300 system is first released with J06.17. The NonStop NS2300 system differs from the NS2100 system as follows: .

- NonStop NS2300 system use a different blade element.
- NonStop NS2300 system has different memory options.

For more information about NonStop NS2300 system, refer to the *NonStop NS2300 Planning Guide*

NonStop NS2200 Series Systems

NonStop NS2200 series systems use the NS2200 CPU, which has two microprocessor cores enabled. NonStop NS2200 series systems have two VIO enclosures, one VIO enclosure for each ServerNet fabric, and connections for processors 0-3 are made through ports 1-4 of the VIO Logic Board in slot 14 of each VIO enclosure. For more information on NonStop NS2200 series systems, refer to the *NonStop NS2200 Series Planning Guide*.

NonStop NS2200 series systems support connections to CLIMs and SAS disk drive enclosures. NonStop NS2200 series systems do not support FCDMs, and they do not support connections to NonStop S-series I/O enclosures.

The method for booting is also different from the method for NS2000 series systems. See [“Starting a System”](#) for more information.

For more information on NonStop NS2200 series systems, refer to the *NonStop NS2200 Series Planning Guide*.

NonStop NS2100 Systems

The difference between NonStop NS2200 series systems and NonStop NS2100 systems is that in the NonStop NS2100 system the NS2100 CPU blade element has one microprocessor core enabled. For more information on NonStop NS2100 systems, refer to the *NonStop NS2100 Planning Guide*.

NonStop NS2000 Series Systems

The difference between NonStop NS2200 series or NonStop NS2100 systems and NonStop NS2000 series systems is that NonStop NS2000 series systems support different blade elements: NonStop NS2000 series systems use NS2000 CPU blade elements, but NonStop NS2200 series systems use NS2200 CPU blade elements, and NonStop NS2100 systems use NS2100 CPU

blade elements. For more information on NonStop NS2000 series systems, refer to the *NonStop NS2000 Series Planning Guide*.

NonStop NS2000 series systems support connections to CLIMs and SAS disk drive enclosures. NonStop NS2000 series systems do not support FCDMs, and they do not support connections to NonStop S-series I/O enclosures.

The method for booting is different from the method for NonStop NS2200 series systems or NonStop NS2100 systems. See “[Starting a System](#)” for more information.

NonStop NS16000 Series Systems

In NonStop NS16000 series 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.

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

NonStop NS16000 series servers running H06.20 or later RVUs support IP, Telco, and Storage CLIMs.

NonStop NS16000 and NS16200 servers running H06.23 or later RVUs support IB CLIMs.

Volume level encryption is supported on NonStop NS16000 series servers that are licensed to support it. If your system is licensed for volume level encryption, you can encrypt disk and tape volumes that are connected to the system with a Storage CLIM. See the *NonStop Cluster I/O Module (CLIM) Installation and Configuration Guide* for details.

NonStop NS14000 Series Systems

In NonStop NS14000 series systems, there are no p-switches. There are now two types of NS14000 series 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 3 through 5 of the two IOAMs in the IOAM enclosure for communications to storage devices and subsystems as well as to LANs). NonStop NS14000 systems do not support connections to additional IOAM enclosures or NonStop S-series I/O enclosures.
- A NonStop NS14000 series 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.

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

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

NonStop NS1000 and NS1200 Systems

NonStop NS1000 and NS1200 systems have no processor switches or LSUs. Like NonStop NS14000 systems, there are now two types of NS1000 systems: those consisting of a single IOAM enclosure and those consisting of one VIO enclosure for each fabric (two VIO enclosures). NS1200 systems consist of VIO enclosures only. ServerNet connectivity for each type is accomplished as described for the “[NonStop NS14000 Series Systems](#)”, except for the absence of the LSUs.

NonStop NS1000 and NS1200 systems do not support connections to NonStop S-series I/O enclosures. Besides the architectural differences, NonStop NS1000 and NS1200 systems also utilize different NonStop Blade Elements than NonStop NS16000 series or NS14000 series systems. For more information on NonStop NS1000 and NS1200 systems, refer to the planning guide for your NonStop system.

NonStop NS1000 and NS1200 systems do not support connections to IP CLIMs, Storage CLIMs, Telco CLIMs, IB CLIMs, or SAS disk drive enclosures.

Terms Used to Describe System Hardware Components

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

- “[Device](#)” (page 38)
- “[System Resource or Object](#)” (page 38)
- “[Server and System](#)” (page 38)
- “[Blade](#)” (page 39)

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 system components that OSM software and other HP manageability tools (for example, HP SIM) display, monitor, and sometimes control. 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).

Server and System

In NonStop NS2400 series, NS2300, NS2200 series, NS2100, NS2000 series, NS16000 series, NS14000 series, NS1200, and NS1000 systems, “server” refers to the hardware, and “system” refers to the combination of the hardware and the software running on it.

In NonStop BladeSystems, “system” can refer collectively to all of the assembled and tested hardware: the HP NonStop BladeSystem c7000 enclosure, the NonStop server blades, the processing subsystem, the networking subsystem, and the storage subsystem. It can also refer to the combination of that hardware and the software running on it.

“NonStop server” refers to a NonStop NS2400 series server, NS2300 server, NS2200 series server, NS2100 server, NS2000 series server, NS16000 series server, NS14000 series server,

NS1200 server, or NS1000 server. Or it can refer to a NonStop server blade in a NonStop BladeSystem, NonStop NS2200 series system, NS2100 system, or NS2000 series system.

Blade

The term “blade” is used differently in different NonStop systems:

- In NonStop BladeSystems, NonStop server blades house the microprocessors and are mounted inside c7000 enclosures.
- In NonStop NS16000 series and NS14000 series systems, a NonStop Blade Element (NSBE) consists of a chassis, processor board containing two to four PEs, memory, I/O interface board, midplane, optics adapters, fans, and power supplies. The OSM Service Connection displays the Blade Complex object. A Blade Complex houses two NSBEs (in a duplex, or DMR system) or three NSBEs (in a triplex, or TMR system) in each enclosure. A Blade Complex consists of up to four logical processors and their associated LSUs.
- In NonStop NS2400 series, NS2300, NS2200 series, NS2100, NS2000 series, NS1200, and NS1000 systems, a NonStop Blade Element (NSBE) is an HP Integrity server adapted for use as a Blade Element. A Blade Element contains the processor elements, power supplies, fan assemblies, and firmware. NonStop NS2400 series, NS2300, NS2200 series, NS2100, and NS2000 series systems can include two or four Blade Elements. NonStop NS1200 and NS1000 systems can include up to eight Blade Elements.

In NonStop NS2400 series, NS2300, NS2200 series, NS2100, and NS2000, NS1200, and NS1000 systems, the OSM Service Connection displays the Blade Complex object, and there is one Blade Element per Blade Complex. The OSM Service Connection uses the Blade Complex to hierarchically differentiate between each Blade Element in NonStop NS2400 series, NS2300, NS2200 series, NS2100, NS2000 series, NS1200, and NS1000 systems.

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 or the OSM System Inventory Tool 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.
The OSM System Inventory Tool allows you to inventory multiple systems in a single action and save the data to a Microsoft Excel file. For more information, see the online help available within the OSM System Inventory Tool.
- 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 appropriate planning guide for your NonStop system.

Using SCF to Determine Your System Configuration

You can use SCF commands to configure and control the objects (lines, controllers, processes, and so on) belonging to each subsystem running on a NonStop NS-series system or NonStop BladeSystem. 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. For complete information, see the *SCF Reference Manual for J-Series and H-Series RVUs*. Subsystem-specific information appears in a separate manual for each subsystem. For a partial list of these manuals, refer to [“Related Reading for Tools and Utilities” \(page 227\)](#).

More details about the functions of SCF appear in [“Subsystem Control Facility \(SCF\)” \(page 225\)](#).

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 J-Series and 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.CONF $xxyy$ 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 CONF $xxyy$, 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” \(page 193\)](#).

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 $LI.#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 1 \(page 42\)](#), the SCF LISTDEV command lists all the physical and logical devices on the system.

Example 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.TSYS DP2
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.TSYS DP2
132	\$DATA5	0,297	1,286	(3,42)	4096	220	\DRP14.\$SYSTEM.SYS00.TSYS DP2
133	\$DATA4	0,298	1,285	(3,44)	4096	220	\DRP14.\$SYSTEM.SYS00.TSYS DP2
134	\$DATA3	0,299	1,284	(3,42)	4096	220	\DRP14.\$SYSTEM.SYS00.TSYS DP2
135	\$DATA2	0,300	1,283	(3,42)	4096	220	\DRP14.\$SYSTEM.SYS00.TSYS DP2
136	\$DATA1	0,301	1,282	(3,44)	4096	220	\DRP14.\$SYSTEM.SYS00.TSYS DP2
137	\$DATA	0,302	1,281	(3,44)	4096	220	\DRP14.\$SYSTEM.SYS00.TSYS DP2
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 1 \(page 42\)](#) 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 1 gives the names of some subsystems that are common to most NonStop NS-series systems and are routinely monitored by operations. These subsystems appear in the LISTDEV output in Example 1 (page 42).

Table 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 1 (page 42) 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 J-Series and 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 NonStop systems. The examples use the SCF ASSUME command to make a given subsystem the current default object for gathering information.

CIP Subsystem

The CIP subsystem supports CLIMs in NonStop BladeSystems and in NonStop NS-series systems that support CLIMs. The CIP subsystem cannot be totally determined via SCF. Instead, use `climconfig` via `CLIMCMD`, as described in “CIP Subsystem Configuration” (page 50).

TCP/IP Subsystem

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

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

Table 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 ¹	INFO SUBNET *
All ROUTE names ¹	INFO ROUTE *

¹ For SUBNET and ROUTE objects in the CIP subsystem, refer to “CIP Subsystem Configuration” (page 50).

NonStop systems 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 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 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 4](#) to display desired information.

Table 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

Table 4 Displaying Information for the Storage Subsystem (\$ZZST0) (continued)

To Display Information About These Configured Objects	Enter This Command
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](#).

Example 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 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 NonStop systems. 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 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 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 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 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 17 (page 99).

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

Table 7 Subsystem Objects Controlled by SCF

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
CIP	Cluster I/O Protocols	48	0 or 1
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

Table 7 Subsystem Objects Controlled by SCF *(continued)*

Subsystem Acronym	Description	Device Type	Device Subtype
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 3](#):

Example 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.QATRACK
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 4](#):

Example 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 5](#):

Example 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 6](#) for Expand-Over-NAM and Expand-Over-ServerNet line-handler processes.

Example 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
```

CIP Subsystem Configuration

The Cluster I/O Protocols (CIP) subsystem provides hardware and a configuration and management interface for I/O on NonStop BladeSystems.

Because CIP provides networking configuration on the CLIM rather than solely on the NonStop host, some of the information that is available for the other TCP/IP subsystems is found on the CLIM. For example, the SCF command line interface for the SUBNET and ROUTE objects in the CIP subsystem only provide information about the host-side objects. To obtain information about the network interface and the network route, you must use the CLIMCMD `climconfig` command line interface.

This TAACL command shows all SUBNET names for the CIP subsystem:

```
>climcmd <clim-identifier> climconfig interface -info all
```

This TAACL command shows all ROUTE names for the CIP subsystem:

```
>climcmd <clim-identifier> climconfig route -info all
```

3 Overview of Monitoring and Recovery

- “When to Use This Chapter” (page 51)
- “HP Tools for Monitoring System Resources” (page 52)
 - “What the HP Tools Monitor” (page 52)
 - “Monitored NonStop Hardware Resources” (page 52)
- “Using HP SIM” (page 55)
- “Using the HP SIM Plug-Ins” (page 57)
 - “NonStop Software Essentials” (page 57)
 - “NonStop Software Essentials” (page 57)
 - “NonStop Cluster Essentials” (page 58)
 - “NonStop Cluster Performance Essentials” (page 58)
 - “Insight Control Power Management” (page 58)
 - “HP Insight Remote Support Advanced” (page 59)
- “Using the OSM Service Connection” (page 59)
 - “Recovery Operations for Problems Detected by OSM” (page 59)
 - “Monitoring Problem Incident Reports” (page 59)
 - “Generating Diagnostic Data” (page 60)
- “Using SCF” (page 60)
 - “Determining Device States” (page 60)
 - “SCF Object States” (page 62)
- “Using Onboard Administrator” (page 63)
- “Automating Routine System Monitoring Procedures” (page 63)
 - “Using Automated HP Manageability Tools” (page 64)
 - “Using TACL Macros, Routines, or Command Files” (page 64)
- “Using the Status LEDs to Monitor the System ” (page 68)
- “Related Reading for NonStop Monitoring” (page 70)

When to Use This Chapter

Use this chapter to learn about the HP manageability tools that operators can use to monitor and manage NonStop hardware resources. It also refers you to other chapters or manuals for more detailed information about monitoring and managing specific NonStop components, events, applications, and processes.

HP Tools for Monitoring System Resources

A number of HP tools, several of them automated and GUI-based, are available from which you can view, and in some cases manage, the status of NonStop hardware resources. They include:

- [“Using HP SIM” \(page 55\)](#)
- [“Using the HP SIM Plug-Ins” \(page 57\)](#)
- [“Using the OSM Service Connection” \(page 59\)](#)
- [“Using SCF” \(page 60\)](#)
- [“Using Onboard Administrator” \(page 63\)](#)

See the chapter [“Monitoring the Performance of NonStop Subsystems” \(page 71\)](#) for a description of the HP tools used to monitor the performance of NonStop subsystems.

What the HP Tools Monitor

The HP manageability tools:

- View alarms for the resources in the system. Look for alarm icons and red and yellow triangle icons.
- View the attributes of a selected resource in the system. If you suspect that a resource is having problems, you can check the resources to identify information such as the name of the resource, type of resource, and the state of the resource.
- If needed, configure the system console to provide dial-out problem information to a service provider or allow a service provider or an authorized customer to dial-in to your server.
- Check for incident reports. For example, the OSM server generates problem incident reports when critical changes occur to a system resource. Problem incident reports are visible in EMS events and HP Instant Remote Support Advanced. See [“Monitoring Problem Incident Reports” \(page 59\)](#).
- Use the manageability tools that monitor NonStop performance.
- Check for event messages. For example, HP SIM and the OSM Event Viewer display events from event log files.
- Check the internal X and Y fabrics to examine the system’s connections between its groups.
- Check for logged status or error messages generated by the HP manageability tool.

Monitored NonStop Hardware Resources

[Table 8 \(page 53\)](#) identifies NonStop system hardware resources that are monitored and the HP tools that monitor them.

NOTE: The OSM Service Connection sends alarm information to HP SIM, allowing HP SIM to display overall system health as well as specific alarms (“alerts” in HP SIM) for affected system resources. Thus, for every system resource in [Table 8 \(page 53\)](#) that lists the OSM Service Connection as a monitoring tool, HP SIM also provides a level of monitoring and notification for those resources. The OSM Service Connection remains the necessary tool for serviceability actions as well as detailed attribute information for those resources. For the system resources in which HP SIM is listed as the monitoring tool, those resources can only be monitored by HP SIM and not the OSM Service Connection.

For a complete and detailed summary and descriptions of all the NonStop system and cluster resources monitored and managed by the OSM Service Connection, including their attributes and the actions that can be performed from the OSM Service Connection, see the *OSM Service Connection User’s Guide*.

For information about the manageability tools you use to monitor NonStop subsystem performance, see [“Monitoring the Performance of NonStop Subsystems”](#) (page 71).

Table 8 Summary of Monitored NonStop Hardware Resources

Resource	Monitored using these tools	For more information, see..
Adapters for communications subsystems: G4SA	OSM Service Connection SCF interface to various subsystems	“Using the OSM Service Connection” (page 59) Chapter 7 (page 85) Chapter 9 (page 106) <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” (page 59) Chapter 9 (page 106) <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>
BladeCluster Solution	OSM Service Connection	<i>BladeCluster Solution Manual</i>
c7000 enclosure and subcomponents, including fans, power supplies, batteries	Onboard Administrator (OA) OSM Service Connection	<i>HP BladeSystem Onboard Administrator User Guide</i> “Using the OSM Service Connection” (page 59) <i>OSM Service Connection User’s Guide</i> (or OSM Service Connection online help)
Communications lines	SCF interface to the various subsystems	Chapter 7 (page 85)
Disk drives (in general and in specific configurations or locations)	OSM Service Connection SCF interface to the storage subsystem DSAP	“Using the OSM Service Connection” (page 59) Chapter 12 (page 140) <i>Guardian User’s Guide</i>
Fibre Channel Disk Module (FCDM) enclosures (and individual disk drives) attached to FCSAs	OSM Service Connection SCF interface to the storage subsystem DSAP	“Using the OSM Service Connection” (page 59) Chapter 9 (page 106) Chapter 12 (page 140) <i>Guardian User’s Guide</i>
Networking CLIMs	CLIMCMD OSM Service Connection	Chapter 10 (page 112) “Using the OSM Service Connection” (page 59) <i>OSM Service Connection User’s Guide</i> (or OSM Service Connection online help)
Modular I/O adapter module (IOAM) and subcomponents, including ServerNet switch boards, power supplies, and fans	OSM Service Connection	“Using the OSM Service Connection” (page 59) “Monitor Batteries” (page 173) <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” (page 59) Chapter 11 (page 120) <i>OSM Service Connection User’s Guide</i> (or OSM Service Connection online help)
NonStop server blade and its subcomponents, including enclosure, ServerNet interface card and DIMMs	Integrated Lights-Out (iLO) Onboard Administrator (OA) OSM Service Connection	<i>HP Integrity BL860c Server Blade User Service Guide</i> <i>HP BladeSystem Onboard Administrator User Guide</i>

Table 8 Summary of Monitored NonStop Hardware Resources *(continued)*

Resource	Monitored using these tools	For more information, see..
NonStop ServerNet Cluster 6770 Switch	OSM Service Connection	“Using the OSM Service Connection” (page 59) <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 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 System Console (NSC)	HP SIM	<i>HP SIM for NonStop Manageability</i>
Power Regulator for NonStop BladeSystems NB54000c, NB54000c-cg, NB56000c, NB56000c-cg	Insight Power Management plug-in for HP SIM OSM	“Insight Control Power Management” (page 58) <i>HP SIM for NonStop Manageability</i> <i>OSM Service Connection User’s Guide</i>
Printers	SCF SPOOLCOM	Chapter 14 (page 162) <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	<i>OSM Service Connection User’s Guide</i> (or OSM Service Connection online help)
SAS disk drive enclosure and subcomponents, including I/O module, power supply, fan, disk	OSM Service Connection	“Using the OSM Service Connection” (page 59) <i>OSM Service Connection User’s Guide</i> (or OSM Service Connection online help) <i>NonStop BladeSystem Planning Guide</i>
ServerNet connectivity for a NonStop BladeSystem or for a NonStop NS 14000 series, NS1200, or NS1000 system (which have no processor switches) 4-Port ServerNet Extender (4PSE) installed in an IOAM enclosure (NS14000 or NS1000 only) or VIO enclosure (NS14000 series, NS1200, or NS1000)	OSM Service Connection	“Using the OSM Service Connection” (page 59) <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” (page 59) <i>OSM Service Connection User’s Guide</i> (or OSM Service Connection online help) Chapter 8 (page 100)

Table 8 Summary of Monitored NonStop Hardware Resources *(continued)*

Resource	Monitored using these tools	For more information, see..
ServerNet switch (in place of processor switch in NonStop BladeSystems)	OSM Service Connection	“Using the OSM Service Connection” (page 59) <i>OSM Service Connection User’s Guide</i> (or OSM Service Connection online help)
ServerNet wide area network (SWAN) concentrator	OSM Service Connection SCF interface to the WAN subsystem	“Using the OSM Service Connection” (page 59) <i>OSM Service Connection User’s Guide</i> (or OSM Service Connection online help) Chapter 7 (page 85)
Storage CLIMs	CLIMCMD OSM Service Connection	Chapter 10 (page 112) “Using the OSM Service Connection” (page 59) <i>OSM Service Connection User’s Guide</i> (or OSM Service Connection online help)
Tape drives	OSM Service Connection SCF interface to the storage subsystem MEDIACOM	Chapter 13 (page 155) Chapter 9 (page 106) <i>Guardian User’s Guide</i>
Uninterruptible Power Supply (UPS)	OSM Service Connection	“Monitor Batteries” (page 173) <i>OSM Service Connection User’s Guide</i> (or OSM Service Connection online help)
Virtual TapeServer (VTS), as well as a Modular Storage Adapter (MSA) attached to a VTS	HP SIM	<i>HP SIM for NonStop Manageability</i>

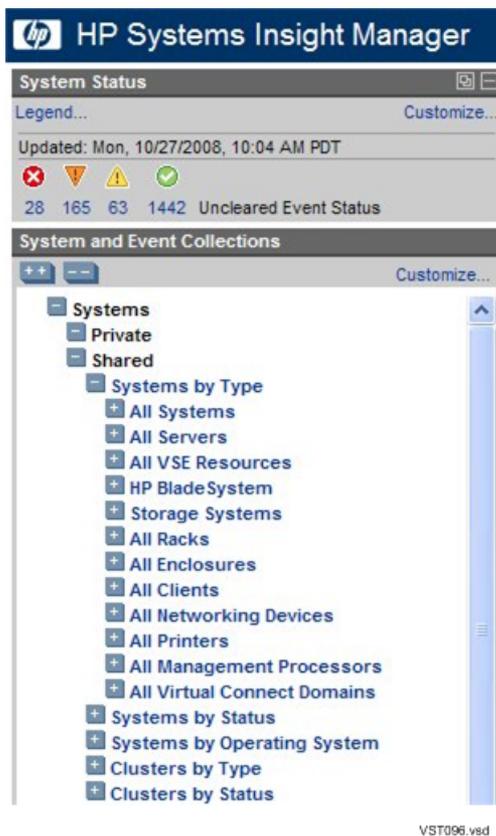
Using HP SIM

HP Systems Insight Manager (HP SIM) is the primary HP solution for using a single interface to manage NonStop servers and storage as well as other HP servers. HP SIM provides a very high-level display of overall system health for user-selected NonStop systems and NonStop system clusters. It does not actively monitor these systems, but instead displays the system health information for them that it receives from the NonStop Open System Management (OSM) software and several of the HP SIM plug-in products.

The HP SIM plug-ins allow users to perform a variety of tasks: install and manage NonStop server software, configure and control CLIMs from the Cluster I/O Protocols (CIP) subsystem, manage clusters of NonStop systems, and monitor and manage temperature and power consumption for NonStop BladeSystems NB54000c, NB54000c-cg, NB56000c, and NB56000c-cg. For more information about the HP SIM plug-ins, see [“Using the HP SIM Plug-Ins” \(page 57\)](#) and *HP SIM for NonStop Manageability*.

From the navigation pane in HP SIM, you can choose to view all or selected systems in a single window, based on criteria such as system status, system type, and so forth. See [Figure 1](#).

Figure 1 HP SIM Tree View



HP SIM includes these general features for all platforms:

- Installs on Windows, HP-UX, and Linux CMS (Central Management Service)
- Inventory, fault, and configuration management
- Security features: role-based authorizations, OS-based authentication, SSL/SSH support
- Distributed task facility to remotely run commands, scripts, and batch files on managed systems

NOTE: Only Windows CMS is currently supported to manage NonStop systems.

In order to use HP SIM to display information about one or more NonStop systems, you must first discover the systems, as described in *HP SIM for NonStop Manageability*.

The HP BladeSystem Integrated Manager is a component in HP SIM that provides streamlined management for the entire HP BladeSystem product family. This plug-in correlates information from Onboard Administrators (OAs), Integrated Lights Out processors (iLOs), and OS-based management agents (OSM server for NonStop) to show the entire data center environment. For example, it can show all the racks, which enclosures are in which rack, and which blades are running what OS. Then, use this view to drill down to that system. HP BladeSystem Integrated Manager installs automatically with HP SIM.

While supported on NonStop system consoles, HP recommends installing HP SIM and SIM plug-ins on a separate Central Management Server (CMS). For more information, including considerations for whether to install those products on system consoles or a separate CMS, see *HP SIM for NonStop Manageability*. You can download HP SIM from <http://www.hp.com/go/hpsim> or install it from the HP Insight Control for NonStop DVD, which ships along with the NonStop System Console Installer DVD (see the *HP NonStop System Console Installer Guide*).

For more about using HP SIM, including the guidelines and requirements for installing HP SIM, Insight Remote Support Advanced, and HP SIM plug-ins, including considerations for Central Management Server (CMS) and NonStop system consoles, see *HP SIM for NonStop Manageability*,

the HP SIM online help, or manuals listed for HP Systems Insight Manager under Network and Systems Management on docs.hp.com.

Using the HP SIM Plug-Ins

These HP SIM plug-in products provide value-add, in some cases necessary, services in the NonStop environment:

- “NonStop I/O Essentials” (page 57)
- “NonStop Cluster Essentials” (page 58)
- “NonStop Cluster Performance Essentials” (page 58)
- “Insight Control Power Management” (page 58)
- “HP Insight Remote Support Advanced” (page 59)

NonStop Software Essentials

The NonStop Software Essentials plug-in installs and manages NonStop server software. NonStop Software Essentials uses DSM/SCM server software and provides a replacement for the DSM/SCM Planner Interface tasks—as well as certain Host Maintenance Interface functions—that is more secure, easier to use, and overcomes other shortcomings of DSM/SCM. NonStop Software Essentials coexists with DSM/SCM; you can use either NonStop Software Essentials or DSM/SCM for tasks in which there is overlapping functionality.

NOTE: The DSM/SCM Planner Interface is not supported on 64-bit system consoles (including NonStop System Consoles running Windows Server 2008), so the NonStop Software Essentials interface must be used on those consoles.

For more information, see the *NonStop Software Essentials Installation and Quick Start Guide*, the *NonStop Software Essentials User’s Guide*, and the online help for NonStop Software Essentials.

NonStop I/O Essentials

The NonStop I/O Essentials plug-in provides support for all configuration and control tasks in the CIP (Cluster I/O Protocols) subsystem for NonStop systems with CLIMs. It provides a graphical user interface alternative to the command line interfaces of the CLIMCMD tool and SCF. It can perform actions on multiple CLIMs in different systems with a single action.

Key benefits of using NonStop I/O Essentials are:

- Ability to perform all CIP management tasks without knowing the syntax of NonStop or Linux commands.
- Ability to perform CIP management tasks on multiple targets (for example CLIMs, disks, tapes, and so on) in parallel.
- Pre-created management tasks provided to configure and control all aspects of CIP.
- Secure and encrypted communication between the system console and NonStop systems. All data in motion between the console and the NonStop system is encrypted over SSH (Secure Shell)
- Remote management via a web browser.

NonStop I/O Essentials can be used stand-alone or with NonStop Cluster Essentials. When used in conjunction with NonStop Cluster Essentials, all the functionality of NonStop I/O Essentials can be deployed across a cluster of NonStop systems.

For more information about installing NonStop I/O Essentials, refer to the *NonStop I/O Essentials Installation and Quick Start Guide* and the online help in NonStop I/O Essentials.

NonStop Cluster Essentials

The NonStop Cluster Essentials plug-in acts as a central integration point for cluster management. It provides integrated cluster management for homogeneous clusters of NonStop systems or Unix systems and for heterogeneous clusters of NonStop systems and Unix systems. It supports clusters consisting of NonStop BladeSystems, NonStop NS-series systems, NonStop S-series servers and optionally, Proliant systems running Linux.

After you discover a NonStop system from HP SIM, NonStop Cluster-specific pages are displayed in HP SIM. The NonStop Clusters home page lists all NonStop clusters that have been configured. The home page for an individual NonStop cluster shows all systems that are in that cluster.

For more information about installing NonStop Cluster Essentials, refer to the *NonStop Cluster Essentials Installation and Quick Start Guide*. Once the NonStop Cluster Essentials plug-in is installed, refer to the online help available from the NonStop Clusters home page for information about using NonStop Cluster Essentials.

NonStop Cluster Performance Essentials

NonStop Cluster Performance Essentials is a plug-in to NonStop Cluster Essentials that monitors homogeneous clusters of NonStop systems and heterogeneous clusters of NonStop and Linux systems. With it, you can monitor the performance of individual NonStop or Linux systems or clusters of them, using the HP SIM and NonStop Cluster Essentials user interfaces. NonStop Cluster Performance Essentials provides these key features:

- Performance monitoring of CPU, Processor, Disk, Transaction, Expand, SQL Statement, ServerNet, File, and CLIM entities
- Alert notifications for defined and configured performance threshold values
- Security by encryption of all communications between NonStop Cluster Performance Essentials and the managed systems

For more information about NonStop Cluster Performance Essentials, refer to the *NonStop Cluster Performance Essentials User Guide* and *HP SIM for NonStop Manageability*.

Insight Control Power Management

Insight Control Power Management is an HP SIM plug-in that displays the layout of servers and devices in racks and data center floors, summarizing temperature and power consumption. It allows you to monitor and manage power capacity and utilization for blades and c-class enclosures and power delivery devices for data centers. It also allows you to monitor inlet air temperature and 24-hour peak temperature for blades and c-class enclosures. Insight Control Power Management provides both power and thermal monitoring for NonStop BladeSystems. It also provides alerts for potential errors, lack of redundancy and configuration anomalies.

OSM has been enhanced to allow a NonStop system to be powered down gracefully from Insight Control Power Management.

For NonStop BladeSystems NB54000c, NB54000c-cg, NB56000c, and NB56000c-cg only, Insight Control Power Management should be used to set the Power Regulator mode immediately after the Power Regulator has been enabled in the OSM Service Connection. See *HP SIM for NonStop Manageability* for procedure information.

The Insight Control Power Management installer is included on the HP Insight Control for NonStop DVD that ships with the NonStop System Console Installer DVD. For information about installing Insight Control Power Management, see the *NonStop System Console Installer Guide*.

For information about using Insight Control Power Management, see the *HP Insight Control Power Management User Guide* for the version you are using.

HP Insight Remote Support Advanced

HP Insight Remote Support Advanced is the go-forward remote support solution for all NonStop systems, replacing the OSM Notification director in both modem-based and HP Instant Support Enterprise Edition (ISEE) remote support solutions. For more information about Insight Remote Support Advanced, refer to the HP Integrity NonStop Service information collection of NTL.

Using the OSM Service Connection

The OSM Service Connection is the primary OSM management component and interface for system monitoring and serviceability.

See [“Overview of OSM Applications” \(page 29\)](#) for examples of how the other OSM applications are used for monitoring-related functions.

- 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.
- In T0682H02 AAM and later, OSM has enhanced automatic data collection capabilities. When enabled, OSM automatically collects diagnostic data whenever a hardware failure causes a Problem Incident Report (IR) to be created. For more information, see the *OSM Service Connection User's Guide* (also available as online help within the OSM Service Connection).

For detailed information about using the OSM Service Connection to monitor and manage NonStop components, see the *OSM Service Connection User's Guide*.

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.
- Examining, or forwarding to your service provider, the information collected by OSM as part of its automatic data collection capabilities (see [“Using the OSM Service Connection” \(page 59\)](#)).
- Check the section in this guide that covers the system resource—for example, [Chapter 13 \(page 155\)](#)— for information on using the SCF and other tools to determine the cause of a problem. Then follow the directions in the Recovery Operations section 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 Collection” \(page 31\)](#).

Monitoring Problem Incident Reports

Problem incident reports are the means by which remote support services are provided for NonStop systems. Although the OSM Notification Director might still be used to provide such services, HP Insight Remote Support Advanced is the go-forward solution for remote support, replacing the OSM Notification Director in both modem-based and HP Instant Support Enterprise Edition (ISEE) remote support solutions.

With Insight Remote Support Advanced, OSM server software still performs problem diagnosis on the NonStop system and creates incident reports (IRs), which are passed along to the NonStop system console or HP Central Management Server (CMS) via WEBES. IRs are displayed as events in HP SIM and, when so configured, are forwarded to HP support.

For more information about Insight Remote Support Advanced, refer to the HP Integrity NonStop Service information collection of NTL.

Generating Diagnostic Data

The OSM Service Connection provides several options for setting up for and generating diagnostic information about the system resources it monitors. Among them is the System action **Collect Diagnostic Data**, which generates Diagnostic Data files for the system. For detailed information about this and other commands that set up for or perform diagnostic data generation, see the *OSM Service Connection User's Guide*.

Using SCF

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 ["Related Reading for Printers"](#) (page 163).

Determining Device States

This section 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 7](#) shows the results of the SCF STATUS TAPE \$* command:

Example 7 SCF STATUS TAPE Command

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

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

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

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

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

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

STORAGE - Status TAPE \COMM.$DLT25
LDev  State      Primary  Backup  DeviceStatus
      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.

```
subsystem STATUS object-type object-name
```

Name	State	PPID	BPID	attr1	attr2	attr3	...
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 9 lists and explains the possible object states that the SCF STATUS command can report.

Table 9 SCF Object States

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.
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.

Table 9 SCF Object States *(continued)*

State	Substate	Explanation
	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.
	UNKNOWN-REASON	The input/output (I/O) process is down for an unknown reason.
STOPPING		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.

Using Onboard Administrator

The Onboard Administrator (OA) is the enclosure management processor, subsystem, and firmware base used to support the c7000 enclosure and the NonStop server blade enclosures in it. To launch the OA from the OSM Service Connection, right-click an enclosure object, select **Actions**, and perform the **Launch OA URL** action. The OA provides additional information on the enclosure, such as details on fans and power supplies in the enclosure.

Redundant OA modules in the c7000 enclosure manage the power and cooling of that enclosure, any other c7000 enclosures in the NonStop BladeSystem, and the NonStop server blades in the c7000 enclosures. The OA software is integrated with OSM and iLO (Intergrated Lights Out) via Single Sign-On (SSO). Only one instance of the OA operates at a time, but switchover is automatic if one fails. For more information about using the OA, see to the *Onboard Administrator User Guide*.

Automating Routine System Monitoring Procedures

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

Using Automated HP Manageability Tools

You should use the automated, GUI-based HP manageability tools listed in [Table 10 \(page 64\)](#) to create automated system monitoring procedures:

Table 10 Management Applications for Automating System Monitoring Procedures

Management application	For more information, see..
SIM Performance Essentials	<i>NonStop Cluster Performance Essentials User Guide</i> <i>HP SIM for NonStop Manageability</i>
Web ViewPoint and its plug-ins	“Web ViewPoint” (page 73) <i>Web ViewPoint User Guide</i>
NonStop Availability Statistics and Performance (ASAP)	“NonStop Availability Statistics and Performance (ASAP)” (page 72) <i>HP NonStop ASAP Server Manual</i> <i>HP NonStop ASAP Client Manual</i> <i>HP NonStop ASAP QuickStart Guide</i>
HP Operations Agent for NonStop (OVNM) and HP Performance Agent for NonStop (OVNPM)	“HP Operations Agent for NonStop (OVNM)” (page 75) <i>HP Operations Agent for NonStop User and Installation Guide for Windows</i> “HP Performance Agent for NonStop (OVNPM)” (page 75) <i>HP Performance Agent for NonStop Server Monitoring Guide</i>

Using TACL Macros, Routines, or Command Files

If you don't have any of the automated management applications listed in [Table 10 \(page 64\)](#) to automate system monitoring procedures, your operations environment can use 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 8](#) 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 8](#) into an EDIT file.

Example 8 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 9](#). This output shows that all elements of the system being monitored are up and running normally.

Example 9 System Monitoring Output File

```

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

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

```

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

                RUNNING  STOPPED  THAWED      FROZEN  FREEZE
                17       0        17          0      PENDING
SERVERCLASSES

                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 system is first powered on. [Table 11](#) lists some of the status light-emitting diodes (LEDs) and their functions.

Table 11 Status LEDs and Their Functions

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 FCSA. 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.

Table 11 Status LEDs and Their Functions *(continued)*

Location	LED Name	Color	Function
	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.
	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.
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.	

Table 11 Status LEDs and Their Functions *(continued)*

Location	LED Name	Color	Function
Locator	Flashing Blue	Lights when the system locator is activated.	
P-switch PICs	Power-on	Green	Lights when power is on with PIC available for normal operation.
		Amber	Lights when a fault exists.
P-switch PIC ServerNet connector	Power-on	Green	Lights when a ServerNet link is functional.

For status LEDs on the NonStop BladeSystem c7000 enclosure, NonStop server blades, CLIMs, and SAS disk drive enclosures, see the *NonStop BladeSystem Planning Guide*, which provides references to the currently supported component manuals.

Related Reading for NonStop Monitoring

For more information about monitoring, see the documentation listed in [Table 12](#).

Table 12 Related Reading for Monitoring Information

Task	Tool	For information, see..
Monitoring system hardware, including failed or failing FRUs	HP SIM and plug-ins OSM Service Connection	<i>HP SIM for NonStop Manageability</i> <i>OSM Service Connection User's Guide</i> (online help or NTL version)
Monitoring NonStop BladeSystem enclosures and NonStop server blade enclosures, including fans and power supplies	Onboard Administrator (OA)	<i>HP BladeSystem Onboard Administrator User Guide</i>
Monitoring CLIMs, CLIM components, and CLIM-attached storage	CLIMCMD SCF OSM	<i>Cluster I/O Protocols (CIP) Configuration and Management Manual</i> OSM Service Connection online help
Perform actions on multiple CLIMs with a single action	NonStop I/O Essentials combined with NonStop Cluster Essentials	<i>NonStop I/O Essentials Installation and Quick Start Guide</i> Online help for NonStop I/O Essentials and NonStop Cluster Essentials
Using SCF, its commands and options, and device types and subtypes	SCF interface to subsystems	<i>SCF Reference Manual for H and J Series RVUs</i> <i>SCF Reference Manual for the Storage Subsystem</i>
Monitoring clustered systems	OSM Service Connection	<i>BladeCluster Solution Manual</i> <i>ServerNet Cluster 6780 Operations Guide</i> <i>ServerNet Cluster Manual</i> OSM Service Connection online help
Obtaining and installing HP SIM	HP SIM	<i>HP SIM for NonStop Manageability</i> <i>NonStop Console Installer Guide</i>

4 Monitoring the Performance of NonStop Subsystems

This chapter provides an overview of the HP manageability tools you can use to monitor the performance of NonStop entities.

- “Summary of NonStop Performance Objects Monitored” (page 71)
- “HP Manageability Tools That Monitor NonStop Performance” (page 72)
 - “NonStop Availability Statistics and Performance (ASAP)” (page 72)
 - “Web ViewPoint NonStop Storage Analyzer Plug-In” (page 74)
 - “HP Operations Agent for NonStop (OVNM)” (page 75)
 - “HP Performance Agent for NonStop (OVNPM)” (page 75)

Summary of NonStop Performance Objects Monitored

Table 13 (page 71) identifies the NonStop performance-monitoring tools and their NonStop objects.

Table 13 NonStop Performance-Monitoring Tools and Their Objects

NonStop Objects Monitored	Performance-Monitoring Tool	For more information, see..
CPU, Process, Disk, Transaction, Expand, SQL Statement, ServerNet, File, CLIM	NonStop Cluster Performance Essentials, a plug-in to HP SIM, monitors these performance objects, which you can access from the HP SIM PE column	“Using HP SIM” (page 55) <i>NonStop Cluster Performance Essentials User Guide</i> <i>HP SIM for NonStop Manageability</i>
Application, CLIM, Comm (Async, SNAX, X25), Processor, Disk, Expand, File, Process, RDF, Spooler, Swap, Tape, TCP/IP, Telserv, TMF	NonStop Availability Statistics and Performance (ASAP)	“NonStop Availability Statistics and Performance (ASAP)” (page 72) <i>HP NonStop ASAP Server Manual</i> <i>HP NonStop ASAP Client Manual</i> <i>HP NonStop ASAP QuickStart Guide</i>
CPU, Disk, Process, File, TMF, Expand	Web ViewPoint	“Web ViewPoint” (page 73) <i>Web ViewPoint User Guide</i>
Disk, File	Web ViewPoint NonStop Storage Analyzer Plug-In	“Web ViewPoint NonStop Storage Analyzer Plug-In” (page 74) <i>Web ViewPoint Storage Analyzer Plug-In User Guide</i>
App, Base24, Cpu, Disk File, Expand Line, Pathway, Process, Program, ServerNet, Socket, Spooler, SQL Procedure, SQL Statement, TCP/IP, TMF	HP Operations Agent for NonStop (OVNM) and HP Performance Agent for NonStop (OVNPM)	“HP Operations Agent for NonStop (OVNM)” (page 75) <i>HP Operations Agent for NonStop User and Installation Guide for Windows</i> “HP Performance Agent for NonStop (OVNPM)” (page 75) <i>HP Performance Agent for NonStop Server Monitoring Guide</i>

HP Manageability Tools That Monitor NonStop Performance

You can use these HP manageability tools to obtain performance information, both current and historical, for NonStop subsystems:

- “NonStop Cluster Performance Essentials” (page 72)
- “NonStop Availability Statistics and Performance (ASAP)” (page 72)
- “Web ViewPoint NonStop Storage Analyzer Plug-In” (page 74)
- “HP Operations Agent for NonStop (OVNM)” (page 75)
- “HP Performance Agent for NonStop (OVNPM)” (page 75)

NonStop Cluster Performance Essentials

The NonStop Cluster Performance Essentials software provides plug-in components to HP SIM to support performance monitoring clusters of NonStop servers or Linux servers, or a combination of the two. The software allows you to monitor the performance of individual clusters or individual servers using the HP SIM interface.

NonStop Cluster Performance Essentials consists of:

- Performance Essentials HP SIM client (PE SIM Plug-In)
- Performance Essentials NonStop Host (NonStop Agent)

The Performance Essentials SIM plug-in to HP SIM runs on the Central Management Server (CMS) and the NonStop agent runs on NonStop servers. Performance Essentials agents collect and route performance metrics from the managed servers to the Performance Essentials SIM plug-in, enabling performance monitoring of the NonStop servers from the CMS using server-specific performance data. Agent software runs on each managed NonStop server in association with and in support of the SIM plug-in component. Together, the PE SIM plug-in and agents gather performance metrics from managed NonStop clusters, representing them in the HP SIM GUI, determining the status of clusters periodically, and providing alert notifications based on user-defined thresholds.

For more information about NonStop Performance Essentials, see the *NonStop Cluster Performance Essentials User Guide*. For more information about HP SIM, see *HP SIM for NonStop Manageability*.

NonStop Availability Statistics and Performance (ASAP)

The NonStop Availability Statistics and Performance (ASAP) software product monitors the availability and performance of NonStop systems and application objects:

- System object information includes the operational and performance status of NonStop system components.
- Application information includes user-defined application domain availability statistics.

You can use NonStop ASAP to obtain real-time and historic views of your NonStop operating system, to see everything from generalized overviews of availability and performance for system and application objects to detailed, customized views of specific node-entity-object-attribute combinations.

NonStop ASAP consists of these monitoring components:

- The ASAP Server includes remote-node statistics-gathering processes (SGPs) that run throughout an Expand network of NonStop systems. These intelligent SGPs gather and analyze Application,

CLIM, Processor, Expand File, Process, RDF and other entity information and report it to a local-node ASAP database server.

- The ASAP Server collects and normalizes network-wide application and system availability data and stores this information in the ASAP database. The ASAP local-node server includes command interpreters and database servers which act on behalf of ASAP workstation clients.
- The ASAP Client includes GUI and intelligent agent functions that allow you to view and analyze system and application statistics. Using ASAP, you can view the overall state of your NonStop systems and applications, as well as performance information. You can focus on specific nodes, entities, and objects.

From the ASAP client, you can open and use these windows and view to obtain NonStop system and application information:

- The Browse window allows you to browse and analyze node-entity-object availability and performance statistics. You can see both data grid and graph views of the specified area. The window updates automatically in real time based on the selected nodes, entities, and objects.
- The OSC window provides a high-level graphical overview of object state changes and a detailed log of state changes. You can obtain both high-level and detailed state change information.
- The Graph window shows information in a compact graphical form useful for maintaining multiple views of nodes, entities, and objects.
- The Report window shows detailed reports about a selected node, entity, or object in a Browse or Graph window. You open Report windows from a Browse or Graph window by using context-sensitive pop-up windows.
- The Session window controls your Host session. You can start and stop Host Sessions using the window's menus or toolbar, and monitor Host Sessions through the window's Session log. The Session log contains information about host communications and software versions.
- The ASAP Properties window allows you to customize which entities and attributes are displayed from the Browse and Graph windows.
- OIL tree view provides a hierarchical view of host nodes, entities, and objects so that you can view the overall availability and performance of many nodes, entities, and objects. Availability and Performance icons are displayed in the OIL tree view to guide you to the nodes, entities, and objects with the highest alerts.

For more information, see the NonStop ASAP documentation, including the *NonStop ASAP Client Manual*, *NonStop ASAP Server Manual*, and *NonStop ASAP Quickstart Guide*.

Web ViewPoint

Web ViewPoint is a browser-based interface that monitors the performance and manages the operations of NonStop subsystems. Information is displayed in a GUI format where alert icons identify components that require attention. A graphical drill-down interface is also provided. In most cases, clicking on a NonStop component's displayed data field exposes other related information as well as the commands applicable to that component.

Web ViewPoint uses several plug-ins to add functional components to the Web ViewPoint menus, including:

- ASAP plug-in
- Storage Analyzer plug-in
- Event Analyzer plug-in

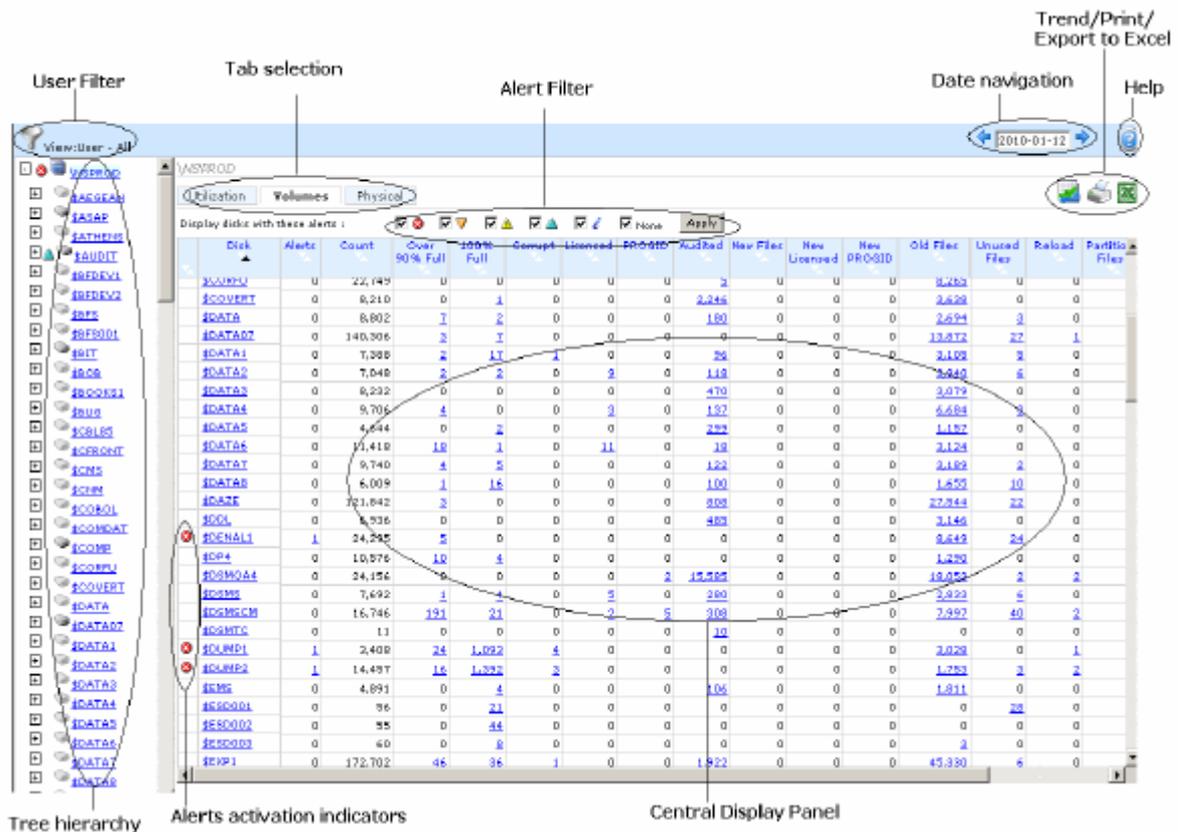
For more detailed information about Web ViewPoint, see the *Web ViewPoint User Guide*.

Web ViewPoint NonStop Storage Analyzer Plug-In

NonStop Storage Analyzer, a Web ViewPoint plug-in, is an automated, web-based software tool you use to analyze the performance of and manage NonStop disk storage resources. It automatically scans an entire NonStop system to provide detailed GUI snapshots of all existing disks and files. This information is organized into several GUI views of both physical and logical disk storage. Based on user-defined thresholds, it generates alerts for exception conditions in EMS event messages that you can forward to various destinations. These alerts serve as an early warning trigger on unacceptable disk or file usage levels or changes in configuration features (for example, file security and licensing).

The Web ViewPoint NonStop Storage Analyzer plug-in's main display screen provides information about total disk utilization, disk utilization by group, and utilization by users in a group. It displays information about individual disk volumes, summaries of physical disk usage, partition files, and fragmented files. You can drill down to a particular disk to obtain information about its files and to perform operations on those files. See [Figure 2 \(page 74\)](#).

Figure 2 Main Display Screen for Web ViewPoint NonStop Storage Analyzer plug-in



The main screen has three tabs: Utilization, Volumes, and Summary. The left side contains a hierarchical format tree with a single resident node at its root.

The Central Display Panel on the right uses grids, lists, and details to display the current selection on the tree. It contains:

- In the top left center, three selection tabs: Utilization, Volumes, and Summary. When the Volumes or Summary tab is selected, you can choose from a series of Alert icons with checkable boxes to select to display only disks which have the selected Alerts on them.
- Above the tree, a User Filter to control the scope of content displayed.
- In the upper right corner of the Central Display Panel is Date Navigation, which allows you to display information for a specified date, and a Trend icon that, when clicked, opens a new Trend Display window from which you can select and display trend values for various metrics.

The Disk Utilization tab initially displays a Utilization Summary, which contains three pie chart diagrams of the system's total disk usage by percentage of used and free space, physical disks utilization by groups, and the total number of files used by each group. Below the charts is a table that summarizes the file and size usage for each group.

The Volumes tab displays information about all the volumes and drives that are present in the system.

The Physical tab displays information about physical disks.

The Web ViewPoint NonStop Disk Analyzer plug-in provides methods for obtaining additional performance information about a selected disk and its subvolume on the current node. From the tree hierarchy, you can drill down to a particular disk and its subvolumes to display Disk Properties for an expanded view of their attributes and the alerts generated by them. Clicking on a disk from the tree changes the Central Display. Two tabs, Properties and Sub Volumes, are displayed. The Properties tab can also be displayed by clicking on the desired disk from the Volumes tab in the Central Display section.

The Properties tab displays information for SCF Configuration, Storage, Content, Disk Alerts, File Alerts, and the subvolume trending window. SCF Configuration specifies whether a disk is physical or not. The Show physical link(s) button lets you launch a window that displays the SCF configuration of a disk. This window displays the pool to which the disk belongs and the logical drives that the pool contains.

The Subvolumes tab displays information for Sub Volume, Alerts information, File Count, Total Size, New File Count, and New Size. Sub Volume identifies the name of the subvolume and provides a clickable link to the Subvolume information for the selected disk.

For more information about the Web ViewPoint NonStop Storage Analyzer plug-in and how to use it, see the *Web ViewPoint WVP NonStop Storage Analyzer Plug-in User Guide*.

HP Operations Agent for NonStop (OVNM)

HP Operations Agent for NonStop (OVNM) integrates HP Operations Software (previously known as HP OpenView) with NonStop systems and provides end-to-end management of your NonStop environment.

You can use OVNM to:

- Manage NonStop system and application events using the OVNM console.
- Define threshold settings on NonStop objects.
- Define service maps of managed objects for the OVNM Console and HP Operation Manager.

For more information, see the *HP Operations Agent for NonStop User and Installation Guide for Windows* and the other OVNM documentation.

HP Performance Agent for NonStop (OVNPM)

HP Performance Agent for NonStop (OVNPM) monitors NonStop system applications and customer applications, as well as NonStop system hardware and system software components.

OVNPM installation creates a default configuration for out-of-the-box use. It scans your NonStop system, automatically builds a set of entities (instances) for the CPU, DISK, GROUP, LINE, NETLINE, NODE, SERVERNET, TMF, and USER domain types, and saves them in the USERCFG configuration file. The OVNPM installation process cannot know or access your user applications, so its entities and the entities for the remaining supported domain types (APPLIDDEF, APPLIFILE, DISKFILE, BASE24DEF, FILESIZE, PATHWAY, PROCESS, PROGRAM, SOCKETS, SQLPROC, SQLSTMT, TCP/IP, and USERDEF) must be added manually by editing the configuration file.

Following installation, you must modify and build the OVNPM configuration whenever these changes are made to your NonStop system:

- Hardware is added or removed.
- System software is added or removed.
- User applications are added or removed.
- A change is made to the monitoring status of an existing OVNPM entity.

To monitor a user applications on your NonStop system, the set of entities (instances) that belong to the OVNPM domains PROCESS, PROGRAM, APPLIFILE, APPLIDDEF, USERDEF, and BASE24DEF must be discovered and added to the USERCFG file. The discovery process involves the steps that build instances for these OVNPM domains:

1. Get the instances related to processes by locating the subvolume or directory that contains the executable and processes that belong to the application .
2. Get the instances related to data files by using the navigator to obtain the data file opened by those processes.
3. Get the USERDEF and BASE24 instances by defining their respective counters.

For information needed to use the HP Performance Agent for NonStop (OVNPM) to monitor your NonStop system, including USERCFG file syntax and modification rules and detailed entity descriptions, instance configurations, and metrics, see the *HP Performance Agent for NonStop Server Monitoring Guide* and other OVNPM documentation.

5 Monitoring EMS Event Messages

- “When to Use This Chapter” (page 77)
- “What Is the Event Management Service (EMS)?” (page 77)
- “Tools for Monitoring EMS Event Messages” (page 77)
- “Related Reading for EMS Event Messages” (page 78)

When to Use This Chapter

Use the chapter 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 NonStop system, use one of these tools:

- “HP SIM Event Status Column and Event Tab” (page 77)
- “OSM Event Viewer ” (page 77)
- “Web ViewPoint” (page 78)
- “EMSDIST” (page 78)

HP SIM Event Status Column and Event Tab

From the HP SIM All Systems screen, the ES (Event Status) column displays aggregate event status with a summary of all uncleared events. Event status is cleared whenever an event is added, updated, or removed.

To view the events for a selected system or subsystem, click on the Events tab from its HP SIM view page.

For more information, see *HP SIM for NonStop Manageability*.

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](#)” (page 30). For details on how to use the OSM Event Viewer, refer to the online help.

Web ViewPoint

Web ViewPoint, a browser-based product, accesses the Event Viewer, Object Manager, and Performance Monitor subsystems. Web ViewPoint, a browser-based product, 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.

For more information, see the *Web ViewPoint User Guide*.

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*.

Related Reading for EMS Event Messages

For more information about monitoring EMS event messages, see the documentation in [Table 14](#).

Table 14 Related Reading for Monitoring EMS Event Messages

Tool	For information, see..
EMSDIST	<i>Guardian User's Guide</i>
HP SIM	<i>HP SIM for NonStop Manageability</i>
OSM Event Viewer	<i>OSM Event Viewer online help</i>
Web ViewPoint	<i>Web ViewPoint User Guide</i>

6 Processes: Monitoring and Recovery

- “When to Use This Chapter” (page 79)
- “Types of Processes” (page 79)
 - “System Processes” (page 79)
 - “CIP Processes” (page 80)
 - “I/O Processes (IOPs)” (page 79)
 - “Generic Processes” (page 80)
- “Monitoring Processes” (page 80)
 - “Monitoring System Processes” (page 81)
 - “Monitoring IOPs” (page 82)
 - “Monitoring CIP Processes” (page 82)
 - “Monitoring Generic Processes” (page 82)
- “Recovery Operations for Processes” (page 84)
- “Related Reading for Processes” (page 84)

When to Use This Chapter

This chapter provides basic information about the different types of processes for NonStop systems. 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 NonStop NS-series systems and NonStop BladeSystems:

- “System Processes” (page 79)
- “I/O Processes (IOPs)” (page 79)
- “CIP Processes” (page 80)
- “Generic Processes” (page 80)

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.

CIP Processes

Cluster I/O Protocols (CIP) processes provide configuration and management interfaces for I/O between CLIMs and NonStop server blades in NonStop BladeSystems. For information about CIP processes, refer to the *Cluster I/O Protocols (CIP) Configuration and Management Manual*.

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 section briefly provides examples of some of the HP 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.

Automated HP Tools That Monitor Processes

NonStop Availability Statistics and Performance (ASAP), HP Operations Agent for NonStop (OVNM), and HP Performance Agent for NonStop (OVNPM) can monitor processes and alert you to exception conditions or other performance problems. See [Table 15 \(page 81\)](#).

Table 15 Automated HP Tools That Monitor Processes

HP tools	Monitored objects	For more information, see..
NonStop Availability Statistics and Performance (ASAP)	Process	“NonStop Availability Statistics and Performance (ASAP)” (page 72) <i>HP NonStop ASAP Server Manual</i> <i>HP NonStop ASAP Client Manual</i> <i>HP NonStop ASAP QuickStart Guide</i>
HP Operations Agent for NonStop (OVNM) and NonStop Performance Agent for NonStop (OVNPM)	Process, Expand Line	“HP Operations Agent for NonStop (OVNM)” (page 75) <i>HP Operations Agent for NonStop User and Installation Guide for Windows</i> “HP Performance Agent for NonStop (OVNPM)” (page 75) <i>HP Performance Agent for NonStop Server Monitoring Guide</i>
HP SIM using I/O Essentials plug-in	CIP	“Using HP SIM” (page 55) “NonStop I/O Essentials” (page 57) <i>HP SIM for NonStop Manageability</i> <i>NonStop I/O Essentials Installation and Quick Start Guide</i>

Monitoring System Processes

You can use TACL to check that the system processes are up and running in the processors as you intended. From a TACL prompt, enter:

```
> STATUS *
```

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

```

$SYSTEM STARTUP 2> status *
Process          Pri PFR %WT Userid  Program file                               Hometerm
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
$ZNUPI          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.COMMGR	\$ZHOME
\$DSMSCM	0,335	220	P	317	255,255	\$SYSTEM.SYS14.TSYS2P2	\$ZHOME
\$DATA2	0,336	220	P	317	255,255	\$SYSTEM.SYS14.TSYS2P2	\$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.TSYS2P2	\$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

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 CIP Processes

For information about monitoring CIP processes, refer to the *Cluster I/O Protocols (CIP) 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 PID	Backup PID	Owner ID
CLCI-TACL	\$CLCI	STOPPED		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	
MSGMON	\$ZIM12	STOPPED		None	None	
MSGMON	\$ZIM13	STOPPED		None	None	
MSGMON	\$ZIM14	STOPPED		None	None	
MSGMON	\$ZIM15	STOPPED		None	None	
OSM-APPSVR	\$ZOSM	STARTED		2 ,292	None	255,255
OSM-CIMOM	\$ZCOM	STARTED		2 ,294	3 ,288	255,255
OSM-CONFLH-RD	\$ZOLHI	STOPPED		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	
QIOMON	\$ZM12	STOPPED		None	None	
QIOMON	\$ZM13	STOPPED		None	None	
QIOMON	\$ZM14	STOPPED		None	None	
QIOMON	\$ZM15	STOPPED		None	None	
RTACL	\$RTACL	STOPPED		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	
TFDSHLP	\$ZTH12	STOPPED		None	None	

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*.

For recovery operations on CIP processes, refer to the *Cluster I/O Protocols (CIP) Configuration and Management Manual*.

Related Reading for Processes

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*.

For information about CIP processes, refer to the *Cluster I/O Protocols (CIP) Configuration and Management Manual*.

7 Communications Subsystems: Monitoring and Recovery

- “When to Use This Chapter” (page 85)
- “Communications Subsystems” (page 85)
 - “Local Area Networks (LANs) and Wide Area Networks (WANs)” (page 85)
 - “CLuster I/O Modules (CLIMs)” (page 87)
- “Monitoring Communications Subsystems and Their Objects” (page 88)
 - “Monitoring the SLSA Subsystem” (page 88)
 - “Monitoring the WAN Subsystem” (page 90)
 - “Monitoring the NonStop TCP/IP Subsystem” (page 93)
 - “Monitoring the CIP Subsystem” (page 94)
 - “Monitoring Line-Handler Process Status” (page 96)
 - “Tracing a Communications Line” (page 98)
 - “Recovery Operations for Communications Subsystems ” (page 98)
 - “Related Reading for Communications Subsystems” (page 99)

When to Use This Chapter

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

Communications Subsystems

The software that provides users of 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 CIP subsystem does not use the 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 CIP 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). 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 NonStop systems are the SLSA subsystem and the WAN subsystem.

The SLSA subsystem supports parallel LAN I/O operations, allowing NonStop NS-series systems and NonStop BladeSystems 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 NonStop systems include:

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

The Storage CLIMs and IP CLIMs can take the place of G4SA, FCSA, and IOAM on NonStop BladeSystems.

A Telco CLIM supports the Message Transfer Part Level 3 User Adaptation layer (M3UA) protocol.

An IB CLIM contains three Ethernet interfaces for connecting to customer data LAN(s), and one additional interface dedicated to the maintenance LAN.

The IB CLIMs are used as part of a Low Latency Solution. The Low Latency Solution architecture provides a high speed and low latency messaging system for stock exchange trading from the incoming trade server to the NonStop operating system. The solution utilizes the third-party Informatica software for messaging and order sequencing which must be installed separately.

The Low Latency Solution also requires a customer-supplied IB switch and Subnet Manager software either installed on the IB switch or running on another server.

NOTE: All CLIMs and other I/O devices connect to a maintenance LAN, which is a dedicated service LAN. For more information about maintenance LANs, ask your service provider to refer to the *NonStop Dedicated Service LAN Installation and Configuration Guide*.

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

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 NonStop system 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 NonStop system 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. SWAN concentrators can also be connected to CLIMs, but CLIMs and CIP do not use SLSA.

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.

Cluster I/O Modules (CLIMs)

CLIMs provide communications interfaces for networking and storage for NonStop BladeSystems and some NonStop NS-series systems. For the list of CLIMs supported, see the planning guide for your NonStop system.

The CIP subsystem provides the hardware and a configuration and management interface for I/O on NonStop server blades. In the CIP subsystem, networking functionality is distributed between the CLIM and the NonStop server blade.

CIP requires a NonStop BladeSystem or a NonStop NS-series system that supports CLIMs, and it requires the J06.04 or later J-series RVU or the H06.16 or later H-series RVU on some NonStop servers.

The NonStop host system component of the CIP subsystem includes an application programmatic interface (API), transport service providers, and a Subsystem Control Facility (SCF) interface for monitoring the subsystem and configuring storage and networking environments. The Open System Services (OSS) and Guardian socket library API provides access to the CIP subsystem for applications running on the NonStop server blade. CIP provides a socket access method (SAM) process on the NonStop BladeSystem for applications to use as a transport service provider.

The CLIM component of the CIP subsystem provides the physical interface to the network and performs the IP protocol processing. The CLIM includes one ServerNet Packetizer PCIe card that interfaces to the ServerNet fabric of the NonStop BladeSystem. The CLIM connects from its management processor interface (eth0) and its Integrated Lights-Out (iLO) interface to a maintenance switch.

The Subsystem Control Facility (SCF), the `climconfig` tool, and iLO provide monitoring and management of the CIP subsystem. See [Chapter 10 \(page 112\)](#) and the *Cluster I/O Protocols (CIP) Configuration and Management Manual* for more information. To use iLO for operator tasks, such as booting the CLIM, creating users, and managing access, refer to the *HP Integrity and HP 9000 Integrated Lights-Out Management Processor Guide*.

The OSM Service Connection provides several options for obtaining diagnostic data about CLIMs. The System action **Collect Diagnostic Data** generates Diagnostic Data files for the system, allowing you to locate and review CLIM diagnostic data. The CLIM action **Reboot** performs a CLIM reboot; you can select the “Yes” default value to generate a dump of CLIM diagnostic data before the CLIM reboots. For detailed information about these options and others, see the *OSM Service Connection User’s Guide*.

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 for Communications Subsystems”](#) (page 99).

This guide provides some information about automated HP tools and basic commands you can use to identify and resolve common problems. The automated HP tools identified in [“Automated HP Tools That Monitor Communications Subsystems”](#) (page 88) play an essential role of automatically alerting you to current or developing problems. SCF also provides a 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”](#) (page 39).

Automated HP Tools That Monitor Communications Subsystems

NonStop Availability Statistics and Performance (ASAP), NonStop I/O Essentials plug-in to HP SIM, and NonStop Cluster Performance Essentials can monitor communication subsystem objects and alert you to exception conditions or performance problems. See [Table 16](#) (page 88).

Table 16 Automated HP Tools That Monitor Communications Subsystems

HP tools	Monitored objects	For more information, see..
NonStop Availability Statistics and Performance (ASAP)	Comm	“NonStop Availability Statistics and Performance (ASAP)” (page 72) <i>HP NonStop ASAP Server Manual</i> <i>HP NonStop ASAP Client Manual</i> <i>HP NonStop ASAP QuickStart Guide</i>
HP SIM using I/O Essentials plug-in	CLIM, FCSA, G4SA, IOAM	“Using HP SIM” (page 55) “NonStop I/O Essentials” (page 57) <i>HP SIM for NonStop Manageability</i> <i>NonStop I/O Essentials Installation and Quick Start Guide</i> and online help

Monitoring the SLSA Subsystem

This section 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

2. 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

3. 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 section 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 section 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.$ZTC0.*
```

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 $ZTC0.*
```

The system displays a listing similar to:

```
1-> STATUS SUBNET $ZTC0.*
```

```
TCPIP Status SUBNET \SYSA.$ZTC0.*
```

Name	Status
#LOOP0	STARTED
#EN1	STARTED

Monitoring the CIP Subsystem

CIP processes run on both the NonStop BladeSystem or NonStop NS-series system and the CLIM, but they are completely managed from the NonStop host. Monitor CIP processes on the NonStop BladeSystem or NonStop NS-series system with SCF commands and CLIMCMD commands.

Use the following SCF commands to monitor CIP processes on the NonStop BladeSystem or NonStop NS-series system and to monitor the connection to the CLIM:

- `STATS CLIM`, to display protocol statistics for individual processes on the CLIM
- `STATS MON`, to display statistics pertaining to the interaction of CIP processes on the NonStop BladeSystem or NonStop NS-series system with applications using the socket interface
- `STATUS MON`, to display statistics pertaining to the interaction of CIP processes on the NonStop BladeSystem or NonStop NS-series system with individual CLIMs
- `STATUS CLIM`, to discover whether a given CLIM is present. If the CLIM is present, this command also discovers whether the CLIM is started or stopped and whether it is currently being traced. The `detail` option also displays information about the physical condition of the CLIM.

This is an example using `STATS CLIM` to get statistics about CLIM2:

Example 10 STATS CLIM

```
-> STATS CLIM $ZZCIP.CLIM2
CIP Stats CLIM \COCOA.$ZZCIP.CLIM2

Sample Time ... 11 Jun 2006, 23:51:49.000
Reset Time .... 09 Jun 2006, 2:28:39.000

CLIMMON STATS
Event Log Entries..... 0
CLIMAGT Failures..... 0
Restart..... 1
CIPSSRV0 Failures..... 0

CLIMAGT STATS
Event Log Entries..... 7
Buffer denials..... 0
IT-API errors..... 4
Last IT-API error code..... 12
Linux errors..... 2
Last Linux errno..... 22
Current bfr bytes in use.... 0
High bfr bytes in use..... 2048
Total msgs sent..... 289299
Total msgs received..... 289293
Current connected Cpus..... 4
Queued Commands..... 2
High connected Cpus..... 4
Failed Commands..... 0

CIPSSRV0 STATS
Event Log Entries..... 0
Buffer denials..... 0
IT-API errors..... 0
Last IT-API error code..... 0
Linux errors..... 0
Last Linux errno..... 0
Current bfr bytes in use.... 1048727
High bfr bytes in use..... 1048727
Total msgs sent..... 69
Total msgs received..... 76
Total bytes sent..... 1080
Total bytes received..... 1024
Current connected Cpus..... 4
Deferred Sends..... 0
Current TCP Listen Sockets.. 0
Current UDP Sockets..... 0
Current TCP Connections..... 0
High connected Cpus..... 4
Failed Commands..... 0
High TCP Listen Sockets..... 0
High UDP Sockets..... 1
High TCP Connections..... 1
```

This is an example using STATS MON:

Example 11 STATS MON

```
-> STATS MON $ZZCIP.ZCM01
CIP Stats MON \COCOA.$ZZCIP.ZCM00

Sample Time ... 11 Jun 2006, 23:55:55.300
Reset Time .... 07 Jun 2006, 16:15:13.781

SOCKET STATS
Total Recv Socket Reqs..... 0
Total Recv Errors..... 0
Total Send Socket Reqs..... 0
Total Send Errors..... 0
Data Bytes Sent..... 0
Data Bytes Received..... 0
Total Connections Out..... 0
Current TCP Listen Sockets.. 0
Current UDP Sockets..... 0
Current TCP Connections..... 0
Total Connections In..... 0
High TCP Listen Sockets..... 0
High UDP Sockets..... 0
High TCP Connections..... 0

SOCKET SEND SIZE HISTOGRAM
Size 1-128..... 0
Size 257-512..... 0
Size 1025-2048..... 0
Size 4097-8192..... 0
Size 12289-16384..... 0
Size 32769 and larger..... 0
Size 129-256..... 0
Size 513-1024..... 0
Size 2049-4096..... 0
Size 8193-12288..... 0
Size 16385-32768..... 0

MON STATS
Total Messages Sent..... 0
Total Messages Received..... 0
Current Connected CLIMs..... 2
Total Rejected Connections.. 129
High Connected CLIMs..... 2
Deferred Sends..... 0
```

For more information about these commands and monitoring the CIP subsystem, see the *Cluster I/O Protocols (CIP) Configuration and Management Manual*.

Monitoring CLIM Status

To monitor the operation of a CLIM and its TCP/IP protocol behavior, use these CLIMCMD commands from a TACL prompt on the NonStop host system that run on the CLIM itself:

- CLIM `psclim` script, to monitor CLIM-specific processes
- CLIM `climstatus` script, to monitor CLIM-specific resources. Invoke it without options to monitor a wide range of CLIM devices, or specify the following options:
 - The `s` option provides information about the ServerNet.
 - The `l` option provides information about the Local Area Network (LAN).
 - The `r` option provides information about the Kernel IP routing table.
 - The `h` option provides information about the CLIM OS file-system disks and disk space.

For more information about these scripts and monitoring CLIMs, see the *Cluster I/O Protocols (CIP) Configuration and Management Manual*.

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 section 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” \(page 62\)](#). 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:

- Monitor exception conditions generated by the HP manageability tools.
- 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 for Communications Subsystems” \(page 99\)](#). If you are unable to start a required process or object, contact your service provider.

Related Reading for Communications Subsystems

For more information about monitoring and performing recovery operations for communications subsystems, see the manuals listed in [Table 17 \(page 99\)](#). 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 17 Related Reading for Communications Lines and Devices

For Information About...	Refer to...
General information about communications subsystems	<i>Introduction to Networking for HP Integrity NonStop Systems</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>
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>BladeCluster Solution 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 and J 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>
Using iLO for booting the CLIM, creating users, and managing access	<i>HP Integrity and HP 9000 Integrated Lights-Out Management Processor Guide</i>

8 ServerNet Resources: Monitoring and Recovery

- “When to Use This Chapter” (page 100)
- “ServerNet Connectivity” (page 100)
- “ServerNet Communications Network” (page 101)
- “Monitoring the Status of the ServerNet Fabrics” (page 102)
 - “Monitoring the ServerNet Fabrics Using OSM” (page 102)
 - “Monitoring the ServerNet Fabrics Using SCF” (page 103)
- “Related Reading for ServerNet Resources” (page 105)

When to Use This Chapter

Use this chapter to learn about the following:

- Monitoring and performing recovery operations for the internal and external ServerNet fabrics
- How and when a NonStop NS-series system can be connected to legacy NonStop S-series I/O enclosures
- ServerNet connectivity

NOTE: NonStop BladeSystems and NonStop NS 16000 series systems support connectivity to NonStop S-series I/O enclosures. NonStop BladeSystems, NonStop NS16000 series systems, and NonStop NS14000 series systems can be part of the same ServerNet cluster as NonStop S-series systems. For more information, see the *ServerNet Cluster Supplement for NonStop NS-Series Servers* or the *ServerNet Cluster Supplement for NonStop BladeSystems*.

Commercial versions of NonStop NS2400 series, NS2300, NS2200 series, NS2100, NS2000 series, NS14000 series, NS1200, and NS1000 systems do not support connectivity to NonStop S-series I/O enclosures. For more information, see “[Differences Between NonStop NS-Series Systems](#)” (page 36).

NonStop NS2400 series, NS2300, NS2200 series, NS2100, NS2000 series, NS1200 and NS1000 systems do not support ServerNet clusters.

NonStop BladeSystems and NonStop NS16000 series systems support the BladeCluster Solution topology. BladeCluster is a type of ServerNet cluster. For hardware and software requirements, refer to the *BladeCluster Solution Manual*.

- The BladeCluster Solution is available for NonStop BladeSystems.
- The BladeCluster Solution is available for NonStop NS16000 series systems.
- 6770 ServerNet clusters and 6780 ServerNet clusters can co-exist with BladeCluster Solutions.

For information about BladeClusters, see the *BladeCluster Solution Manual*.

ServerNet Connectivity

Table 18 summarizes ServerNet connectivity and system I/O connectivity in NonStop systems. It also lists the relevant planning guides, which supply figures and more information about connectivity.

Table 18 ServerNet and System I/O Connectivity

System	ServerNet Connectivity	System I/O Connectivity	Reference Guides
NonStop BladeSystem	ServerNet switch on c7000 enclosure, IOAMs, 6770	IP CLIM, Storage CLIM with SAS disk drive enclosures,	<i>NonStop BladeSystem Planning Guide</i>

Table 18 ServerNet and System I/O Connectivity *(continued)*

System	ServerNet Connectivity	System I/O Connectivity	Reference Guides
	ServerNet clusters, 6780 ServerNet clusters, BladeClusters	IB CLIM, Telco CLIM, high-speed Ethernet, Fibre Channel, ESS, FCDM, FC Tape, FC SCSI Tape, additional IOAMEs	<i>BladeCluster Solution Manual</i>
NonStop NS16000 series systems	LSUs, p-switches, IOAMs, , 6770 ServerNet clusters, 6780 ServerNet clusters, BladeClusters	IP CLIM, Storage CLIM with SAS disk drive enclosures, IB CLIM, Telco CLIM, high-speed Ethernet, Fibre Channel, ESS, FCDM, FC Tape, FC SCSI Tape, additional IOAMEs, NonStop S-series I/O enclosures.	<i>NonStop NS16000 Series Planning Guide</i> <i>BladeCluster Solution Manual</i>
NonStop NS14000 series system with IOAME	LSUs, IOAMs, 6770 ServerNet clusters, 6780 ServerNet clusters	High-speed Ethernet, Fibre Channel, ESS, FCDM, FC tape, FC SCSI tape, additional IOAMEs	<i>NonStop NS14000 Series Planning Guide</i>
NonStop NS14000 series system with two VIOs	LSUs, VIOs, 6770 ServerNet clusters, 6780 ServerNet clusters	High-speed Ethernet, Fibre Channel, ESS, FCDM, FC tape, FC SCSI tape	<i>NonStop NS14000 Series Planning Guide</i> <i>Versatile I/O (VIO) Manual</i>
NonStop NS2400 series, NS2300, NS2200 series, NS2100, and NS2200 series system (has two VIOs)	VIO	IP CLIM, Storage CLIM with SAS disk drive enclosures, Telco CLIM, high-speed Ethernet, Fibre Channel, ESS, FC Tape, FC SCSI Tape	<i>NonStop NS2400 Series Planning Guide</i> <i>NonStop NS2300 Planning Guide</i> <i>NonStop NS2200 Series Planning Guide</i> <i>NonStop NS2100 Planning Guide</i> <i>NonStop NS2000 Series Planning Guide</i> <i>Versatile I/O (VIO) Manual</i>
NonStop NS1200 system (has two VIOs)	VIOs	High-speed Ethernet, Fibre Channel, ESS, FCDM, FC tape, FC SCSI tape	<i>NonStop NS1200 Planning Guide</i> <i>Versatile I/O (VIO) Manual</i>
NonStop NS1000 system with two VIOs	VIOs	High-speed Ethernet, Fibre Channel, ESS, FCDM, FC tape, FC SCSI tape	<i>NonStop NS1000 Planning Guide</i> <i>Versatile I/O (VIO) Manual</i>
NonStop NS1000 system with IOAME	p-switches, IOAMs	High-speed Ethernet, Fibre Channel, ESS, FCDM, FC tape, FC SCSI tape	<i>NonStop NS1000 Planning Guide</i>

ServerNet Communications Network

The ServerNet communications network is a high-speed network within an 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 and I/O devices. These dual paths can be used simultaneously to improve performance, and to ensure that no single failure disrupts communications among the remaining system components.

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, such as 6770 ServerNet clusters, 6780 ServerNet clusters, and BladeClusters) 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 (such as a 6770 ServerNet cluster, 6780 ServerNet cluster, or BladeCluster), 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* and the *BladeCluster Solution Manual* provide 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 ServerNet fabric-related objects:
 - a. Internal ServerNet X and Y Fabric objects — for NonStop NS-series systems, they are located under the Fabric Group object, but for NonStop BladeSystems they are located at the top level under the System object.
 - b. ServerNet Switches — located under Enclosure 100 (and, optionally, Enclosure 101), the switch in module 2, slot 5 is the X fabric switch for each enclosure, and the switch in module 3, slot 7 is the Y fabric switch.
 - c. External ServerNet X and Y Fabric objects — located under the ServerNet Cluster object (if your system is part of a ServerNet cluster).

3. Check these objects for:
 - a. If an 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 an 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 an object icon is covered by a yellow arrow, there is a problem on a subcomponent of that object. Expand the object to locate the subcomponent object reporting the alarm or problem attribute.

For NonStop NS16000 series systems, under each internal fabric object, you'll find its associated processor switch (p-switch) module and subcomponents.

For NonStop NS2400 series, NS2300, NS2200 series, NS2100, NS2000 series, NS14000 series, NS1200, and NS1000 systems, there are no p-switches. Instead, ServerNet connectivity is provided through 4PSEs located in the VIO enclosures (or IOAMs, where supported).

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
```

```
1-> status servernet $zsnet
```

```
NONSTOP KERNEL - Status SERVERNET
```

```
X-FABRIC
```

TO	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
FROM																
00	UP	UP	UNA													
01	UP	UP	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													
01	UP	UP	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” \(page 104\)](#).
- 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” \(page 104\)](#).

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” \(page 152\)](#).

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” \(page 130\)](#).

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 ServerNet Resources

For more information about the interconnections between NonStop systems and NonStop S-series systems, see the *NonStop NS 16000 Series Planning Guide*, or *NonStop NS 14000 Series Planning Guide*.

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

9 I/O Adapters and Modules: Monitoring and Recovery

- “When to Use This Chapter” (page 106)
- “I/O Adapters and Modules” (page 106)
 - “Fibre Channel ServerNet Adapter (FCSA)” (page 107)
 - “Gigabit Ethernet 4-Port Adapter (G4SA)” (page 107)
 - “4-Port ServerNet Extender (4PSE)” (page 108)
- “Monitoring I/O Adapters and Modules” (page 108)
 - “Monitoring the FCSAs” (page 108)
 - “Monitoring the G4SAs” (page 109)
 - “Monitoring the 4PSEs” (page 110)
- “Recovery Operations for I/O Adapters and Modules” (page 110)
- “Related Reading for I/O Adapters and Modules” (page 111)

When to Use This Chapter

Use this chapter 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 NS-series documentation.

NOTE: Some NonStop systems are shipped with VIO enclosures instead of an IOAM enclosure. VIO enclosures provide the same functionality as IOAM enclosures. If your NonStop system is equipped with VIO enclosures, refer to the planning guide for your system for information about its VIO 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 as follows:

- Fibre Channel ports are in slots 1 and optionally in 7 and are identified as FCSA in OSM.
- Gigabit Ethernet ports are in slots 6 and optionally in 7 and are identified as G4SA in OSM.

For more information on VIO enclosures, see the *Versatile I/O (VIO) Manual*.

I/O Adapters and Modules

Beginning with 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 NonStop I/O peripherals.

Even though the hardware architecture differs from one series of NonStop systems to another, the 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 Fibre Channel Disk Module (FCDM) enclosure that supports fibre channel disks and an Enterprise Storage System (ESS).

Any connection between a NonStop system and a FCDM 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 FCDM 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 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 J-series RVU, refer to the *J06.nn Release Version Update Compendium*. 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 Compendium*.

Gigabit Ethernet 4-Port Adapter (G4SA)

The M8800 Gigabit Ethernet 4-port ServerNet adapter (G4SA) provides Gigabit connectivity between NonStop systems and Ethernet LANs. G4SAs are installed in slots 1 through 5 of an I/O adapter module (IOAM) (except in 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 NonStop NS14000 and NS1000 systems only, 4PSEs provide ServerNet connectivity between processors and the IOAM enclosure (functionality provided by p-switches in a NonStop NS16000 series system). 4PSEs are located in slot one (and optionally slot 2) of each IOAM. They are connected to the processors through LSUs in NonStop NS14000 systems, directly to the processors (with no LSUs) in NonStop NS1000 systems. FCSAs and G4SAs can be installed in slots 3 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” (page 110).

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 an overview of using the OSM Service Connection to monitor system components, refer to “Using the OSM Service Connection” (page 59).

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 19 lists the possible states for the FCSA.

Table 19 Service, Flash Firmware, Flash Boot Firmware, Device, and Enabled States for the FCSA

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.
Device State: Initializing	Processing is starting up.

Table 19 Service, Flash Firmware, Flash Boot Firmware, Device, and Enabled States for the FCSA
(continued)

State	Description
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 an overview of using the OSM Service Connection to monitor system components, refer to [“Using the OSM Service Connection” \(page 59\)](#).

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 20](#) lists the possible states for the G4SA.

Table 20 Service, Device, and Enabled States for the G4SA

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.

Table 20 Service, Device, and Enabled States for the G4SA *(continued)*

State	Description
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.
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 an overview of using the OSM Service Connection to monitor system components, refer to [“Using the OSM Service Connection”](#) (page 59).

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 for I/O Adapters and Modules](#)” (page 111). If you are unable to start a required process or object, contact your service provider.

Related Reading for I/O Adapters and Modules

For more information about monitoring and performing recovery operations for the I/O adapters and the SLSA and Storage subsystems, see the manuals in the following table. The appropriate manual to use depends on how your system is configured.

Table 21 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 Series Planning Guide</i>

10 Cluster I/O Modules (CLIMs): Monitoring and Recovery

- “When To Use This Section”
- “Cluster I/O Modules (CLIMs)”
- “Monitoring the CLIMs with HP SIM” (page 112)
- “Monitoring the CLIMs with SCF” (page 112)
 - “CLIM Device States” (page 114)
 - “Using SCF to Monitor a CLIM with \$ZZCIP” (page 115)
- “Recovery Operations for CLIMs” (page 118)
- “Related Reading for CLIMs” (page 119)

When To Use This Section

Use this section to monitor Cluster I/O Modules and to perform recovery operations.

NOTE: If you need to power off a CLIM, refer to the *Cluster I/O Protocols (CIP) Configuration and Management Manual*.

Cluster I/O Modules (CLIMs)

Cluster I/O Modules (CLIMs) are servers that provide networking and storage protocols.

CLIM operations are managed with NonStop management tools and applications, including OSM, SCF, CLIMCMD, and HP SIM. CLIMs use PCIe cards for ServerNet connection. For security, CLIMs do not allow any customer code or applications, and there is no modem-based dial-out. In addition, only users authorized for the NonStop host system are allowed to interface with the CLIM because a secure shell relationship is established at installation time between the CLIM and the NonStop host system.

There are different kinds of CLIMs supporting different I/O protocols. See the planning guide for your NonStop system for a list of supported CLIMs.

Monitoring the CLIMs with HP SIM

HP SIM monitors CLIMs as part of its high-level monitoring of NonStop systems. For more detailed diagnostic and servicing functionality, use the OSM Service Connection. NonStop I/O Essentials, NonStop Cluster Essentials, and NonStop Cluster Performance Essentials plug-ins extend HP SIM management of CLIMs across NonStop clusters. For information about using HP SIM and its plug-ins see “Using HP SIM” (page 55) and the online help in each of those plug-in products.

Monitoring the CLIMs with OSM

Use the Subsystem Control Facility (SCF), the CLIMCMD command line interface, and the OSM Service Connection to monitor the CLIMs.

For an overview of using the OSM Service Connection to monitor system components, refer to “Using the OSM Service Connection” (page 59).

Monitoring the CLIMs with SCF

To monitor the CLIM with SCF, use the SCF INFO and SCF STATUS commands.

For example, use STATUS CLIM to determine which CLIMs are present and in what state:

Example 12 STATUS CLIM Summary

```
-> STATUS CLIM $ZZCIP.*  
CIP Status CLIM \MYSYS.$ZZCIP.*
```

Name	Present	State	Trace
CLIM1	Yes	STARTED	OFF
CLIM2	Yes	STARTED	1, 2
CLIM3	Yes	STARTED	2

“Monitoring the CIP Subsystem” (page 94) shows examples of other SCF commands. “Monitoring CLIM Status” (page 96) shows the use of scripts that run on the CLIM itself to monitor the operation of a CLIM and its IP protocol behavior. The *Cluster I/O Protocols (CIP) Configuration and Management Manual* provides reference details and examples for using the SCF INFO and SCF STATUS commands to monitor CLIMs.

In OSM, individual CLIM objects are located under a CLIMs container object, which has no attributes (see Figure 3). If the CLIMs container object shows a bell-shaped alarm icon or a yellow triangular icon, indicating a problem within, expand the CLIMs container object and select the individual CLIM object indicating a problem to check it for alarms and/or attributes, such as the Device State attribute discussed in “CLIM Device States” (page 114). Figure 3 shows the CLIMs container object expanded to show the CLIM objects in it and the attributes of one CLIM object.

Figure 3 CLIM Objects in OSM

The screenshot shows the OSM Service Connection interface. On the left is a tree view of the system hierarchy, including BladeCluster, System \BLOSM3, Internal ServerNet X Fabric, Internal ServerNet Y Fabric, CLIM Attached Disks, CLIMs, Enclosure 100, Monitored Service LAN Device, and SAS Disk Enclosures. The selected CLIM object is CLIM \$ZZCIP.C1002531. On the right, the 'attributes' tab is active, displaying a list of properties for this CLIM. The properties are grouped into sections: Physical, X2 Fabric Information, Y2 Fabric Information, Y1 Fabric Information, Physical Location, X1 Fabric Information, Logical, ROM, and iLO. Each property has a corresponding value, such as Model: ProLiant DL385 G2, Part Number: 414109-B21, and Connection Status: Correctly Connected.

Property	Value
Physical	
Model	ProLiant DL385 G2
Part Number	414109-B21
Serial Number	USE740NBFX
X2 Fabric Information	
Configured Location	Fiber (100.2.5.3.2)
Y2 Fabric Information	
Configured Location	Fiber (100.3.7.3.2)
Y1 Fabric Information	
Configured Location	Fiber (100.3.7.3.1)
Physical Location	
Rack Name	
Rack Offset	
X1 Fabric Information	
Configured Location	Fiber (100.2.5.3.1)
Logical	
CIP Software Version	T0853H01_01MAY2009_18MAR2009_AAL
Configured Connections per ServerNet Fabric	2
Configured Operation Mode	Storage
Connection Status	Correctly Connected
Device State	Started
Down Physical Ethernet Interfaces	None
IP Software Version	T0691H01_01MAY2009_AAL_CLIM_U05
Linux Version	Linux 2.6.18-6-clim-amd64 #1 SMP PREEMP
Maintenance Interface IPv4 Address	192.168.37.31
Storage Software Version	T0830H01_01MAY2009_18MAR2009_AAD
Temperature Status	Normal
Up Physical Ethernet Interfaces	eth0
X1 Fabric Connection State	Correctly Connected
X2 Fabric Connection State	Correctly Connected
Y1 Fabric Connection State	Correctly Connected
Y2 Fabric Connection State	Correctly Connected
iLO IPv4 Address	192.168.42.249
ROM	
Compare State	Same
Default File Name	\$\$SYSTEM.SYS00.rom0848
Default File Version	05.29.2008
Version	05.29.2008
iLO	
Compare State	Down-rev
Default File Name	\$\$SYSTEM.SYS00.fw0848
Default File Version	1.70
Version	1.50

VST200.vsd

CLIM Device States

When monitoring CLIMs using the OSM Service Connection, the state of a CLIM should indicate normal operation. Table 22 lists the possible states for a CLIM.

Table 22 CLIM Device States

State	Substate	Explanation
STARTED		The object is logically accessible to user processes.
STARTING		The object is being initialized and is in transition to the STARTED state.
STOPPED		The object is configured improperly.

If there are problems with CLIM subcomponents, they have their own Device State attributes. Refer to the OSM Service Connection online help for more information

Using SCF to Monitor a CLIM with \$ZZCIP

This is an example of using SCF to monitor the detailed status of an IP CLIM on a NonStop BladeSystem.

Example 13 STATUS CLIM Detailed for IP CLIM on NonStop BladeSystem

```
-> status clim $zzcip.C1002583, detail
CIP Detailed Status CLIM \BLITUG.$ZZCIP.C1002583

Mode..... IP
CLIM HW Connection Status. Connected
State..... STARTED
ConnPts..... 2
X1 Location..... Group 100 , Module 2 , Slot 5 , Port 8 , Fiber 3
Expected Y1 Location..... Group 100 , Module 3 , Slot 7 , Port 8 , Fiber 3
X2 Location..... Group 100 , Module 2 , Slot 5 , Port 8 , Fiber 4
Expected Y2 Location..... Group 100 , Module 3 , Slot 7 , Port 8 , Fiber 4
X1 Connection Status..... Connected
Y1 Connection Status..... Connected
X2 Connection Status..... Connected
Y2 Connection Status..... Connected
Trace Status..... OFF

Last Restart Time..... 30 Sep 2008, 18:14:56.000
CLIM Hostname..... C1002583
CIP SW Version..... T0853H01_01NOV2008_17SEP2008_AAB

Network SW Version..... T0691H01_01NOV2008_AAF_CLIM
More text? ([Y],N) y
Storage SW Version.....
Number of Socket Servers.. 1
Linux Version:
Linux version 2.6.18-6-clim-amd64 (Debian 2.6.18.dfsg.1-22etch2hpde1.5)
(support@hp.com) (gcc version 4.1.2 20061115 (prerelease)
(Debian 4.1.1-21)) #1 SMP PREEMPT Fri Sep 19 19:13:09 UTC 2008

Fabric Status:
CIPMON      Mgmt  Data
ZCM00      XY-2  XY-2
ZCM01      XY-1  XY-1
ZCM04      XY-1  XY-1
ZCM05      XY-2  XY-2

CLIM Maintenance Interface Status & IP Addresses:
eth0          Sts:UP Lkp:UP   IPv4: 16.107.146.135
              IPv6: fe80::21c:c4ff:fedd:a002

Maintenance Provider Interface Status & IP Addresses:
eth0:0        Sts:UP Lkp:UP   IPv4: 16.107.146.50

Data Interface Status & IP Addresses:
Flg  Name      Status  LkP  Master / IP Family & Address
eth5  UP          UP          IPv4: 20.20.20.24
              IPv6: fe80::21c:c4ff:fe48:4e
eth4  UP          UP          IPv4: 20.20.20.23
              IPv6: fe80::21c:c4ff:fe48:4f
eth3  UP          UP          IPv4: 20.20.20.22
              IPv6: fe80::21c:c4ff:fe48:4c
eth2  UP          UP          IPv4: 20.20.20.21
              IPv6: fe80::21c:c4ff:fe48:4d
eth1  UP          UP          IPv4: 20.20.20.20
              IPv6: fe80::21c:c4ff:fedd:a000
lo    UP          --          IPv4: 127.0.0.1
              IPv6: ::1
```

```
Interface Failover Configuration and Status:
Flg  Name      Failover/(Assoc) Fovr  Current      Index
```

C1002583.eth5	--	C1002583.eth5	0x0000000003
C1002583.eth4	--	C1002583.eth4	0x0000000004
C1002583.eth3	--	C1002583.eth3	0x0000000005
C1002583.eth2	--	C1002583.eth2	0x0000000006
C1002583.eth1	--	C1002583.eth1	0x0000000007

This is an example of using SCF to monitor the detailed status of a Storage CLIM on a NonStop NS-series system.

Example 14 STATUS CLIM Detailed for Storage CLIM on NonStop NS-Series System

```
-> STATUS CLIM $ZZCIP.* , DETAIL
CIP Detailed Status CLIM \YOSQA14.$ZZCIP.C100261

Mode..... STORAGE
CLIM HW Connection Status. Connected
State..... STARTED
ConnPts..... 2
X1 Location..... Group 100 , Module 2 , Slot 6 , Port 1
Expected Y1 Location..... Group 100 , Module 3 , Slot 6 , Port 1
X2 Location..... Group 100 , Module 2 , Slot 6 , Port 2
Expected Y2 Location..... Group 100 , Module 3 , Slot 6 , Port 2
X1 Connection Status..... Connected
Y1 Connection Status..... Connected
X2 Connection Status..... Connected
Y2 Connection Status..... Connected
Trace Status..... OFF

Last Restart Time..... 01 Oct 2008, 7:59:20.000
CLIM Hostname..... C100261
CIP SW Version..... T0853H01_01NOV2008_24SEP2008_AAB

Network SW Version..... T0691H01_01NOV2008_AAF_CLIM
Storage SW Version..... T0830H01_01NOV2008_17SEP2008_AAB
Number of Socket Servers.. 1
Linux Version:
Linux version 2.6.18-6-clim-amd64 (Debian 2.6.18.dfsg.1-22etch2hpde1.1)
(support@hp.com) (gcc version 4.1.2 20061115 (prerelease)
(Debian 4.1.1-21)) #1 SMP PREEMPT Wed Sep 10 20:27:40 UTC 2008

Fabric Status:
CIPMON      Mgmt  Data
ZCM00      XY-2  XY-2
ZCM01      XY-1  XY-1
ZCM02      XY-1  XY-1
ZCM03      XY-2  XY-2

CLIM Maintenance Interface Status & IP Addresses:
eth0          Sts:UP  Lkp:UP   IPv4: 16.107.174.86
              IPv6: fe80::21e:bff:fed8:f952

Data Interface Status & IP Addresses:
Flg  Name      Status  LkP   Master / IP Family & Address
  lo          UP      --
              IPv4: 127.0.0.1
              IPv6: ::1
```

Recovery Operations for CLIMs

- Examine the contents of the event message log for the CIP subsystem.
- For information on starting, stopping, monitoring, and managing the CIP subsystem, refer to the *Cluster I/O Protocols (CIP) Configuration and Management Manual*.
- If you are unable to start a required process or object, contact your service provider.

Related Reading for CLIMs

For more information about monitoring and managing CLIMs, see the following:.

- *Cluster I/O Protocols (CIP) Configuration and Management Manual*
- *NonStop Cluster Performance Essentials User's Guide*
- OSM Service Connection online help
- NonStop I/O Essentials online help

11 Processors and Components: Monitoring and Recovery

- “When to Use This Chapter” (page 120)
- “Overview of Processors” (page 120)
- “Monitoring and Maintaining Processors” (page 124)
 - “Monitoring Processors Automatically Using TFDS” (page 125)
 - “Monitoring Processor Status Using the OSM Low-Level Link” (page 125)
 - “Monitoring Processor Status Using the OSM Service Connection” (page 125)
 - “Monitoring Processor Performance Using ViewSys” (page 128)
- “Identifying Processor Problems” (page 128)
 - “Processor or System Hangs” (page 128)
 - “Processor Halts” (page 128)
 - “OSM Alarms and Attribute Values” (page 129)
- “Recovery Operations for Processors” (page 129)
 - “Recovery Operations for a Processor Halt” (page 130)
 - “Halting One or More Processors” (page 130)
 - “Reloading a Single Processor on a Running system” (page 130)
 - “Recovery Operations for a System Hang” (page 133)
 - “Enabling/Disabling Processor and System Freeze” (page 134)
 - “Freezing the System and Freeze-Enabled Processors” (page 134)
 - “Dumping a Processor to Disk” (page 134)
 - “Backing Up a Processor Dump to Tape” (page 137)
 - “Replacing Processor Memory” (page 137)
 - “Replacing the Processor Board and Processor Entity” (page 137)
 - “Submitting Information to Your Service Provider” (page 138)
- “Related Reading for Processors” (page 139)

When to Use This Chapter

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

Overview of Processors

Processors in a NonStop system are one of the following, which are described below:

- “Processors in NonStop Systems Running J-Series RVUs” (page 121)
- “NonStop NS16000 Series and NS14000 Series Blade Complex Processors” (page 122)
- “NonStop NS1200 and NS1000 Blade Element Processors” (page 124)

Processors in NonStop Systems Running J-Series RVUs

NonStop BladeSystem server blades and NonStop NS2400 series, NS2300, NS2200 series, NS2100, and NS2000 series blade elements utilize multicore microprocessors (see [Table 23 \(page 121\)](#)). The set of cores on a NonStop server blade or blade element is considered to be one CPU and is configured as one logical processor. Applications do not see the individual cores. There can be two to sixteen logical CPUs (NonStop server blades) per NonStop BladeSystem, or two to four logical CPUs (NonStop server blade elements) per NonStop NS2400 series, NS2300, NS2200 series, NS2100, or NS2000 series system. There are no logical synchronization units (LSUs) in a NonStop BladeSystem or NonStop NS2400 series, NS2300, NS2200 series, NS2100, or NS2000 series system.

All input and output to and from each NonStop server blade is passed via the 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.

In NonStop BladeSystems and NonStop NS2400 series, NS2300, NS2200 series, NS2100, and NS2000 series systems, ServerNet switches control ServerNet communications, rather than the p-switches found in NonStop NS16000 series systems. In a NonStop BladeSystem, the ServerNet switches are in the c7000 enclosure. In a NonStop NS2400 series, NS2300, NS2200 series, NS2100, and NS2000 series system, the ServerNet switches are in the VIOs. There is one ServerNet switch for each of the ServerNet fabrics.

In summary, these terms describe NonStop BladeSystem or NonStop NS2400 series, NS2300, NS2200 series, NS2100, or NS2000 series system processors:

Table 23 Terms, Systems, and Planning Guides

Term, System, Planning Guide	Description
IPU	Instruction Processing Unit. A single core of the microprocessor.
Logical processor	Also known as a CPU. A set of cores sharing the same memory. All IPUs in a logical processor share the same memory map (except a small per-IPU anchor area). There is one X and one Y ServerNet interface per logical processor.
NonStop server blade	Server blade that houses one multicore microprocessor.
NonStop BladeSystem NB56000c or NB56000c-cg server blade See the <i>NonStop BladeSystem Planning Guide</i>	Two to 16 NonStop server blades each utilize a four-core microprocessor and are populated inside NonStop BladeSystem c7000 enclosures. <ul style="list-style-type: none"> Four-core license file allows use of four cores. Two-core license file allows use of two cores. (2 to 16 logical processors per system)
NonStop BladeSystem NB54000c or NB54000c-cg server blade See the <i>NonStop BladeSystem Planning Guide</i>	Two to 16 NonStop server blades each utilize a four-core microprocessor and are populated inside NonStop BladeSystem c7000 enclosures. <ul style="list-style-type: none"> Four-core license file allows use of four cores. Two-core license file allows use of two cores. (2 to 16 logical processors per system)
NonStop BladeSystem NB50000c or NB50000c-cg server blade See the <i>NonStop BladeSystem Planning Guide</i>	Two to 16 NonStop server blades each utilize a two-core microprocessor and are populated inside NonStop BladeSystem c7000 enclosures. <ul style="list-style-type: none"> Two-core license file allows use of two cores. (2 to 16 logical processors per system)

Table 23 Terms, Systems, and Planning Guides *(continued)*

Term, System, Planning Guide	Description
NonStop blade element	Blade element that houses one dual-core microprocessor.
NonStop NS2400 series blade element See the <i>NonStop NS2400 Planning Guide</i>	Two or four NonStop blade elements each utilize a two-core microprocessor and are installed in modular cabinets. (2 to 4 logical processors per system)
NonStop NS2300 series blade element See the <i>NonStop NS2300 Planning Guide</i>	NonStop blade elements each utilize a two-core microprocessor with one core enabled and are installed in modular cabinets. (2 to 4 logical processors per system)
NonStop NS2200 series blade element See the <i>NonStop NS2200 Series Planning Guide</i>	Two or four NonStop blade elements each utilize a two-core microprocessor with one core enabled and are installed in modular cabinets. (2 to 4 logical processors per system)
NonStop NS2100 blade element See the <i>NonStop NS2100 Planning Guide</i>	Two or four NonStop blade elements each utilize a microprocessor with one core enabled and are installed in modular cabinets. (2 to 4 logical processors per system)
NonStop NS2000 series blade element See the <i>NonStop NS2000 Series Planning Guide</i>	Two or four NonStop blade elements each utilize a two-core microprocessor and are installed in modular cabinets. (2 to 4 logical processors per system)

For NonStop BladeSystem NB54000c, NB54000c-cg, NB56000c, or NB56000c-cg, see the *NonStop BladeSystem Planning Guide* to determine the number of cores enabled in a processor or to increase the number of cores enabled in a processor.

NonStop NS 16000 Series and NS 14000 Series Blade Complex Processors

A NonStop NS 16000 series or NS 14000 series NonStop Blade Complex 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. The 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 comprises the dual modular redundant (DMR) NonStop Blade Complex, which is also referred to as a duplex system.

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).

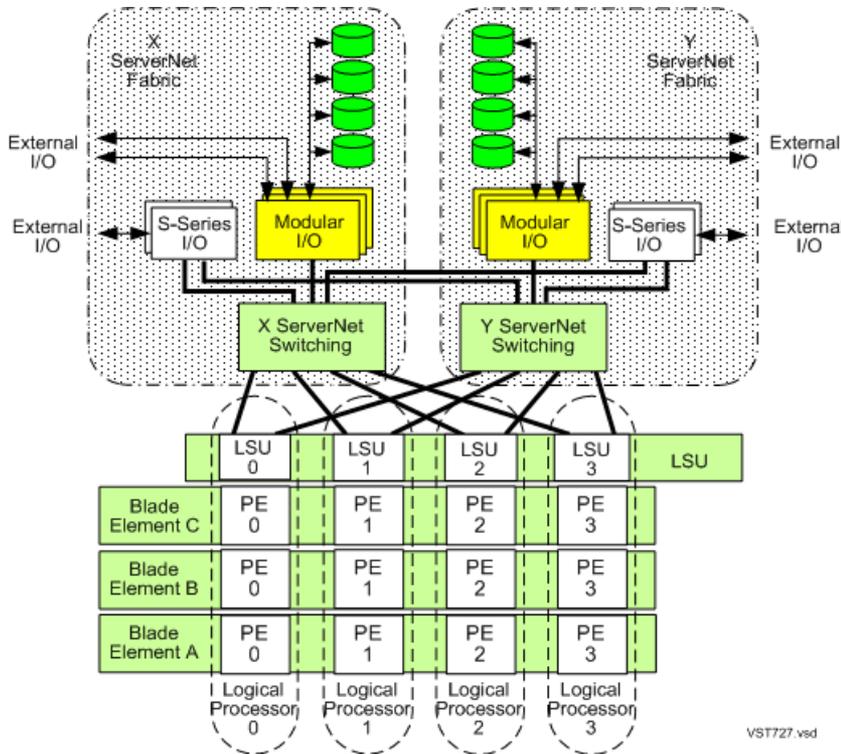
In the event of a processor fault in either a duplex or triplex system, the failed component within a NonStop Blade Element (processor element, power supply, and so forth) can be replaced while

the system continues to run. A single 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 4 shows one NonStop Blade Complex with four processors, the I/O hardware and the ServerNet fabrics.

Figure 4 One NonStop Blade Complex and Four Processors



For NonStop NS16000 series systems, 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 4). 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).

NonStop NS14000 series systems do not have p-switches and cannot be connected to legacy NonStop S-series I/O enclosures. For more information, see [“ServerNet Connectivity” \(page 100\)](#).

In summary, these terms describe NonStop NS16000 series and NS14000 series processors:

Systems and Planning Guides	Term	Description
NonStop NS16000 series systems (2 to 16 processors per system) <i>NonStop NS16000 Series Planning Guide</i> ;	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

Systems and Planning Guides	Term	Description
NonStop NS14000 series systems (2 to 8 processors per system) <i>NonStop NS14000 Series Planning Guide</i>	Blade Element	their associated LSUs. A system includes up to four Blade Complexes. 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.

NonStop NS1200 and NS1000 Blade Element Processors

The NonStop NS1200 or NS1000 blade element is a server blade that contains a single microprocessor and is configured as one logical processor. NonStop NS1200 and NS1000 systems can include up to eight blade elements. There are no logical synchronization units (LSUs) in a NonStop NS1200 or NS1000 system. For more information on NonStop NS1200 and NS1000 systems, see [“Differences Between NonStop NS-Series Systems” \(page 36\)](#) and the appropriate Planning Guide.

All input and output to and from each NonStop blade element is passed via the 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.

In NonStop NS1200 or NS1000 systems, ServerNet switches control ServerNet communications, rather than the p-switches found in NonStop NS16000 series systems. In NonStop NS1200 systems, the ServerNet switches are in the VIOs. In a NonStop NS1000 system, the ServerNet switches are in the VIOs or IOAM enclosure. There is one ServerNet switch for each of the ServerNet fabrics.

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” \(page 125\)](#)
- [“Monitoring Processor Status Using the OSM Low-Level Link” \(page 125\)](#)
- [“Monitoring Processor Status Using the OSM Service Connection” \(page 125\)](#)
- [“Monitoring Processor Performance Using ViewSys” \(page 128\)](#)
- [Chapter 5 \(page 77\)](#)

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 before 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.

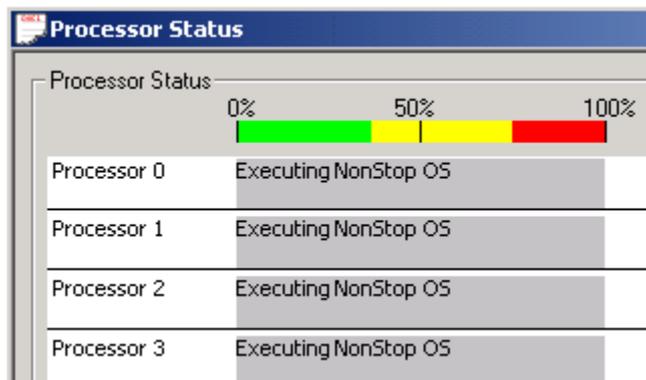
For more information on configuring and using TFDS, see the *Tandem Failure Data System (TFDS) Manual*.

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 5](#).) If not, refer to “[Identifying Processor Problems](#)” (page 128).

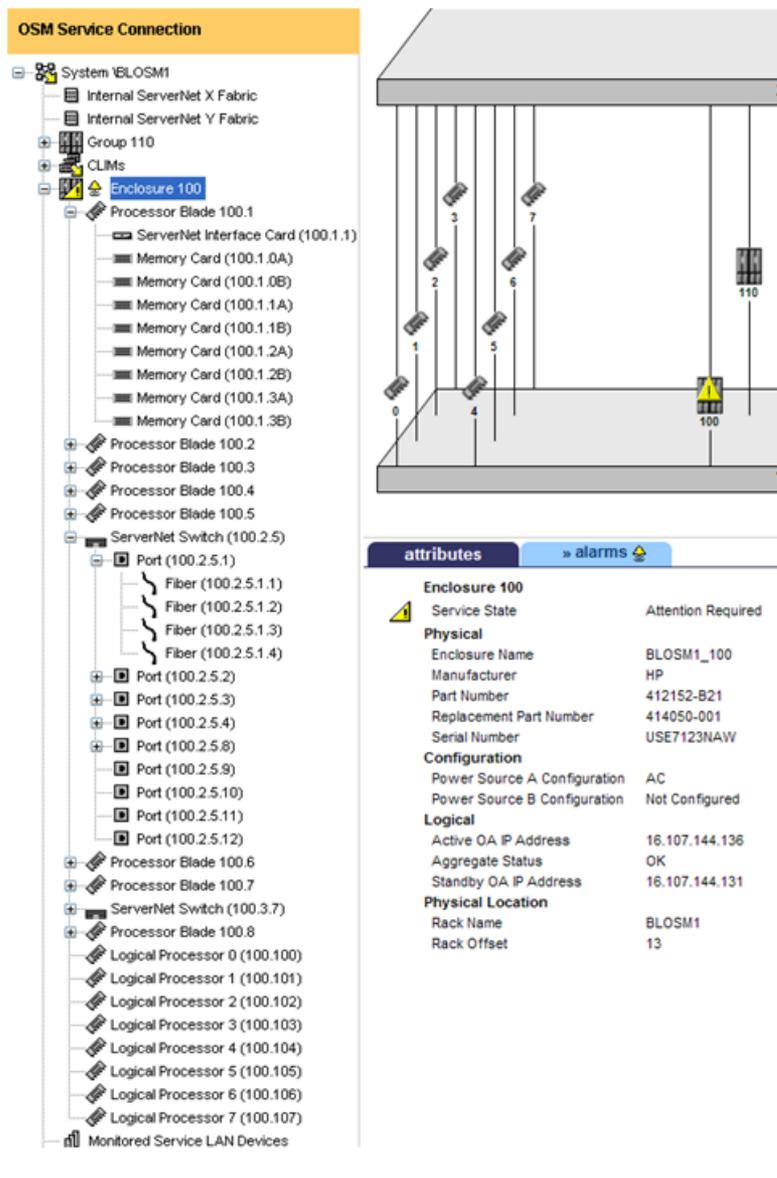
Figure 5 Processor Status Display



Monitoring Processor Status Using the OSM Service Connection

For NonStop BladeSystems, the OSM Service Connection displays processor-related components under the Enclosure object (see [Figure 6](#)), which contains the Processor Blade objects, Logical Processor objects, and ServerNet switch objects, in the tree pane. There can be up to 16 Processor Blade objects per NonStop BladeSystem, representing the total of 16 processors. The Processor Blade objects each contain the Memory Card and ServerNet interface card subcomponents. Each ServerNet switch object contains Port subcomponents, which contain Fiber objects.

Figure 6 OSM Representation of Processor Blade



VST909.vsd

For NonStop NB54000c, NB54000c-cg, NB56000c, and NB56000c-cg BladeSystem server blades, the OSM Service Connection displays the Physical Cores attribute and the Number of Enabled Cores attribute of the Logical Processor object. To increase the number of cores enabled in a processor, obtain a new core license for the new number of cores, and install and enable the new core license according to the *Install Core License* guided procedure on the System object. For instructions to obtain a core license file, refer to the *NonStop BladeSystem Planning Guide*.

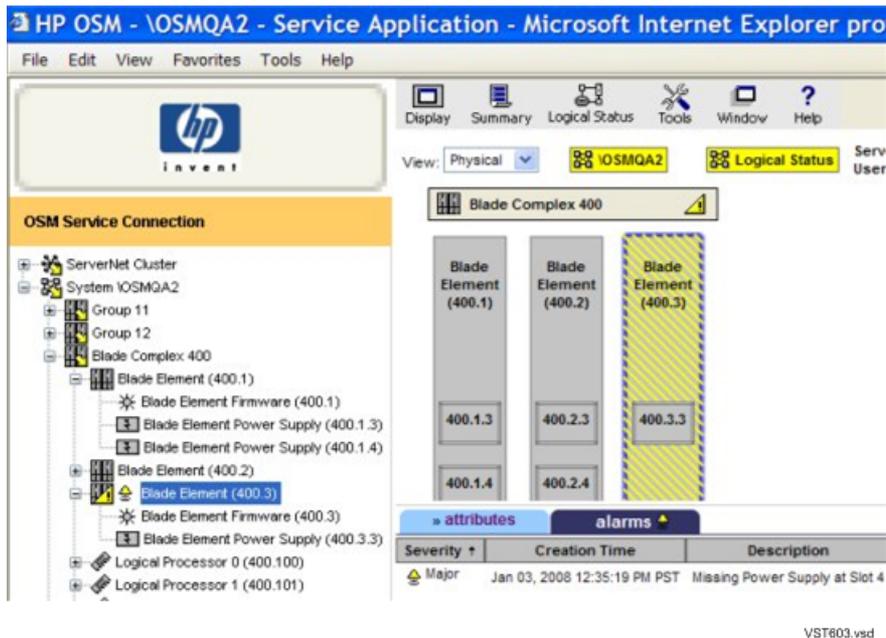
For NonStop NS2400 series, NS2300, NS2200 series, NS2100, NS2000 series, NS16000 series, NonStop NS14000 series, NonStop NS1200, and NonStop NS1000 systems, the OSM Service Connection displays processor-related components under Blade Complex objects in the tree pane.

- A NonStop NS2400 series, NS2300, NS2200 series, NS2100, or NS2000 series system can have two or four Blade Complex objects. When you expand a Blade Complex object (see Figure 7), you should see one Blade Element object and its associated Logical Processor object.
- A NonStop NS16000 series system can have up to four Blade Complex objects with a total of 16 processors. When you expand a Blade Complex object (see Figure 7), you should see up to three Blade Element objects and either two or four Logical Processor objects.

- A NonStop NS14000 series system can have up to three Blade Complex objects with a total of 8 processors. When you expand a Blade Complex object (see [Figure 7](#)), you should see up to three Blade Element objects and either two or four Logical Processor objects.
- A NonStop NS1200 or NS1000 system can have up to eight Blade Complex objects. When you expand a Blade Complex object (see [Figure 7](#)), you should see one Blade Element object and its associated Logical Processor object.

OSM displays blade firmware information on either the Blade Element object or on a Blade Element Firmware object under each Blade Element object, depending on the type of system. The subcomponents, actions, and attributes displayed under Blade Element and Logical Processor objects also depend on the type of system. See the *OSM Service Connection User's Guide* (also available as online help within the application) for more information on what OSM displays for your system type.

Figure 7 OSM Representation of Blade Complex



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

1. Expand the tree pane to check all Blade Complex objects.
2. If a Blade Complex object icon contains a yellow arrow (as illustrated in [Figure 7](#)), 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*.

Managing and Setting NonStop BladeSystems NB54000c, NB54000c-cg, NB56000c, and NB56000c-cg Processor Power Usage with Insight Control Power Management

Power Regulator is available for NonStop BladeSystems NB54000c, NB54000c-cg, NB56000c, and NB56000c-cg only. Power Regulator must first be enabled in the OSM Service Connection before you can set the Power Regulator mode by using Insight Control Power Management. See

HP SIM for NonStop Manageability for these procedures and the supported Power Regulator modes.

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” (page 226).

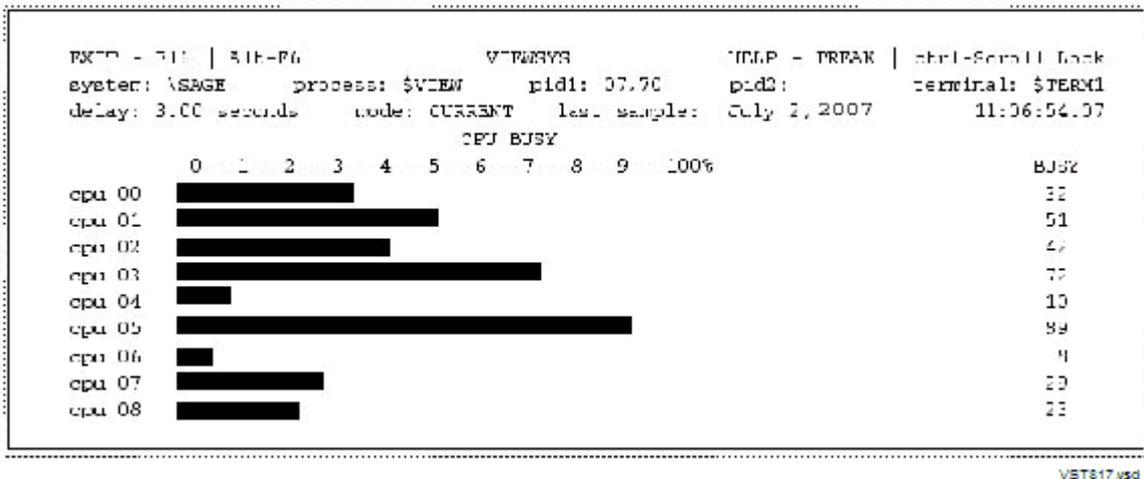
To use ViewSys to obtain information about processor activity, at a TACL prompt:

```
> 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”](#) (page 129).
- Contacting your service provider.

Recovery Operations for Processors

NOTE: All references to processor elements (PEs) and the RELOAD command OMITSLICE parameter apply only to NonStop NS16000 series and NS14000 series systems. For more information on NonStop BladeSystems, NonStop NS2400 series, NS2300, NS2200 series, NS2100 and NS2000 series, NS1200, and NS1000 systems, see [“Differences Between NonStop NS-Series Systems”](#) (page 36) and the appropriate planning guide.

NonStop NS16000 series and NS14000 series systems offer 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.

Processor recovery operations for your system might include:

- [“Recovery Operations for a Processor Halt”](#) (page 130)
- [“Halting One or More Processors”](#) (page 130)

- [“Reloading a Single Processor on a Running system” \(page 130\)](#)
- [“Recovery Operations for a System Hang” \(page 133\)](#)
- [“Enabling/Disabling Processor and System Freeze” \(page 134\)](#)
- [“Freezing the System and Freeze-Enabled Processors” \(page 134\)](#)
- [“Dumping a Processor to Disk” \(page 134\)](#)
- [“Backing Up a Processor Dump to Tape” \(page 137\)](#)
- [“Replacing Processor Memory” \(page 137\)](#)
- [“Replacing the Processor Board and Processor Entity” \(page 137\)](#)
- [“Submitting Information to Your Service Provider” \(page 138\)](#)

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” \(page 125\)](#)). 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” \(page 187\)](#). 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” \(page 134\)](#).

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” \(page 128\)](#)), 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” \(page 134\)](#) and [“Reloading a Single Processor on a Running system”](#) 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 system

Sometimes one or more processors in a running system are not operating. For information on how to determine whether a processor is operating, see [“Monitoring Processor Performance Using ViewSys” \(page 128\)](#).

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” \(page 130\)](#)).

Collect information about the reason for the halt (as described in [“Identifying Processor Problems” \(page 128\)](#)) 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 system are:

- [“Using TACL RELOAD to Perform Reload” \(page 131\)](#)
- [“Using the OSM Service Connection to Perform Reload” \(page 132\)](#)

Following the reload, if TFDS is not configured to take the dump automatically on NonStop NS16000 series and NS14000 series systems, 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 Using RCVDUMP for NonStop NS16000 Series and NS14000 Series Systems” \(page 136\)](#).

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.

For NonStop NS16000 series and NS14000 series systems, 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 OMITSLICE 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, OMITSLICE selects a Blade Element to be excluded. The OMITSLICE parameter applies only to NonStop NS16000 series and NS14000 series systems.

1. From the OSM Low-Level Link, select **File>Start Terminal Emulator**.
2. From the menu, select **For Startup TACL**. If you are using OSM Low-Level Link version T0633AAT or later, an MR-Win6530 session is opened.
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 TACL 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 OMITSLICE parameter. If you don't know which Blade Element to specify, use OMITSLICE 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* OMITSLICE [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.

OMITSLICE [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.

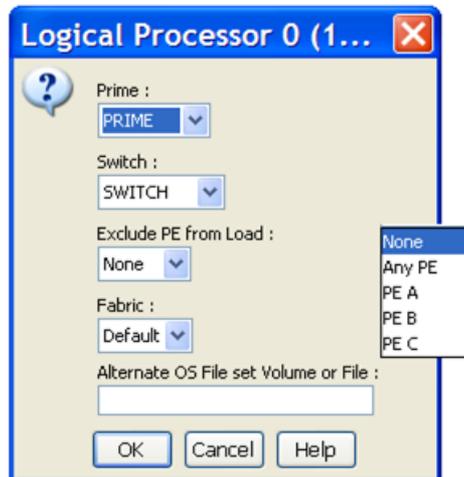
8. Check the TACL window for status messages, which will report successes or errors during the load. Monitor the state of the processor you are loading until it is executing the NonStop Kernel operating system.
9. If the load fails, check the parameters and reload the processor. If the load fails again, contact your service provider.

Using the OSM Service Connection to Perform Reload

For NonStop NS16000 series and NS14000 series systems, 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. The OMITSLICE parameter applies only to NonStop NS16000 series and NS14000 series systems.

- 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**.



VST014.vsd

- 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 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 a system, 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”](#) (page 134)).

2. “Freezing the System and Freeze-Enabled Processors” (page 134).
3. Start the system by loading Processor 0 or 1, as described in “Performing a System Load From a Specific Processor” (page 187). 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” (page 134).
4. Disabling System Freeze (see “Enabling/Disabling Processor and System Freeze” (page 134)).

Enabling/Disabling Processor and System Freeze

NOTE: 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

A processor dump can be made to as many as two files: the full dump file and the extracted dump file. The file name of the full dump file can end with F.

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 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 perform the dump as described in [“Using RCVDUMP to Dump a Processor to Disk”](#) (page 135).

Loading a Down System to Perform a Processor Dump to Disk

To perform a processor dump to disk from a down system, load the system from CPU 0 or CPU 1 with the CIIN file disabled, thereby preventing other processors configured to reload from doing so and allowing you to dump their memory to disk. For detailed instructions, see [“System Load to a Specific Processor”](#) (page 183) and [“Performing a System Load From a Specific Processor”](#) (page 187).

Using RCVDUMP to Dump a Processor to Disk

Using RCVDUMP to dump a processor to disk involves several preparatory steps, followed by performing the RCVDUMP command with the proper syntax.

Before You Begin

- You must have a running CPU and a TACL process. The processor in which the TACL process 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.
2. To empty an existing *dumpfile*:

```
> FUP PURGEDATA dumpfile
```

Dumping a Processor Using RCVDUMP (Excludes NonStop NS16000 Series and NS14000 Series Systems)

NOTE: See [“Dumping a Processor Using RCVDUMP for NonStop NS16000 Series and NS14000 Series Systems”](#) (page 136) for the procedure for NS16000 series and NS14000 series systems.

To use RCVDUMP to dump a processor 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:

NOTE: If you include the `SLICE` or the `PARALLEL` parameter in a RCVDUMP command to any NonStop system other than a NonStop NS16000 series or NS14000 series system, a RCVDUMP error is generated to report: `parameter is not supported on Non-NSAA.`

```
RCVDUMP filename, cpuNum [, SLICE bladeId]
[, START startAddress][, END endAddress]
[, ONLINE | PARALLEL][, FULL]
```

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.

SLICE 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*.

FULL

suppresses creation of an extracted dump file and results in creation of only a full dump file.

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:

Dump completed.

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.

A processor can be dumped to as many as two files: the full dump file and the extracted dump file. The file name of the full dump file may end with F.

Dumping a Processor Using RCVDUMP for NonStop NS 16000 Series and NS 14000 Series Systems

- 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 *SLICE ALL* parameter option.
 - Do not specify *ONLINE* or *PARALLEL* parameters.

See [“Using RCVDUMP to Dump a Processor to Disk” \(page 135\)](#).

- 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 system” \(page 130\)](#)), then use the RCVDUMP command as follows:
 - If more than one Blade Element is in the Stopped state, use the SLICE parameter and specify the *Blade Element (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 SLICE parameter.)
 - Specify the PARALLEL parameter.

See [“Using RCVDUMP to Dump a Processor to Disk” \(page 135\)](#).

- If a dump is to be taken following a system load, as described in [“Performing a System Load From a Specific Processor” \(page 187\)](#), 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 SLICE parameter (RCVDUMP will choose the first running PE).

See [“Using RCVDUMP to Dump a Processor to Disk” \(page 135\)](#).

For NonStop NS16000 series and NS14000 series systems, a processor dump can be made to as many as four files:

- The file name of the full dump file may end with F.
- File names may end with an A, B, or C to designate which Blade Element the processor element was dumped from.

Backing Up a Processor Dump to Tape

NOTE: H-series and J-series NonStop systems do not support backing up a processor dump to tape.

Replacing Processor Memory

Processor memory is field-replaceable for all 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” \(page 128\)](#)), 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 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” (page 138)
- “Web ViewPoint” (page 78)
- “Additional Information Required by Your Service Provider ” (page 138)

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 the following table on one backup tape using the BACKUP utility. Contact your service provider for information about any other files they might need.

Table 24 Other Files to Submit to Your Service Provider

File	Description
\$SYSTEM.ZSYSCONF.CONFIG	System configuration database
\$SYSTEM.SYS _n .CONFTEXT	System configuration file
\$SYSTEM.SYS _n .CONFLIST	System generation program output file
\$SYSTEM.ZLOG _n	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 the following table to your service provider.

Table 25 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.	_____

Table 25 Additional Processor Dump Information for Your Service Provider *(continued)*

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 Processors

For more information about tools used to monitor and perform recovery operations on processors, refer to the documentation listed in the following table.

Table 26 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>
Monitoring processor performance	ViewSys	<i>ViewSys User's Guide</i>
Recovery operations for processor halt	TFDS	<i>Tandem Failure Data System (TFDS) Manual</i>
Service and Support Procedures		Support and Service Collection for NonStop NS-series systems or NonStop BladeSystems

12 Disk Drives: Monitoring and Recovery

- “When to Use This Chapter” (page 140)
- “Overview of Disk Drives ” (page 140)
 - “Internal SCSI Disk Drives” (page 141)
 - “M8xxx Fibre Channel Disk Drives” (page 141)
 - “Enterprise Storage System (ESS) Disks” (page 142)
 - “Serial Attached SCSI (SAS) Disks and Solid State Drives” (page 142)
- “Monitoring Disk Drives” (page 143)
 - “Monitoring Disk Drives With OSM” (page 143)
 - “Monitoring Drives With SCF ” (page 145)
 - “Monitoring the State of Disk Drives” (page 148)
 - “Monitoring the Use of Space on a Disk Volume” (page 149)
 - “Monitoring the Size of Database Files” (page 149)
 - “Monitoring Disk Configuration and Performance” (page 150)
- “Identifying Disk Drive Problems” (page 150)
- “Recovery Operations for Disk Drives ” (page 151)
 - “Recovery Operations for a Down Disk or Down Disk Path” (page 152)
 - “Recovery Operations for a Nearly Full Database File” (page 153)
- “Related Reading for Disk Drives” (page 154)

When to Use This Chapter

Use this chapter to monitor M8xxx Fibre Channel and serial attached SCSI (SAS) drives, and to recover from common disk problems.

Overview of Disk Drives

The NonStop BladeSystems support these types of disk drives:

- “M8xxx Fibre Channel Disk Drives”
- “Enterprise Storage System (ESS) Disks”
- “Serial Attached SCSI (SAS) Disks and Solid State Drives” (page 142)

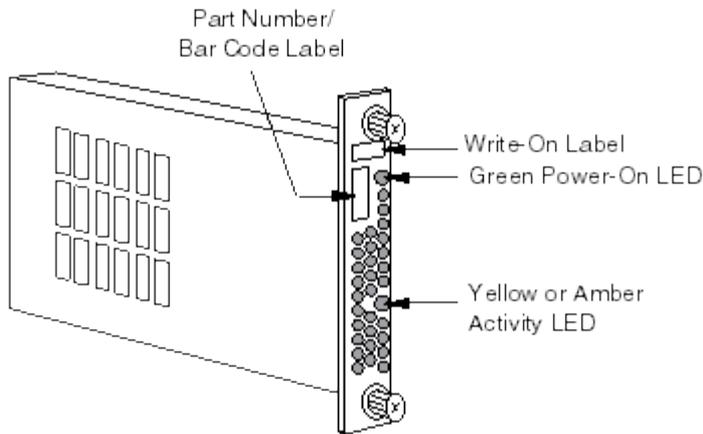
Most NonStop NS-series systems support the types of disk drives listed below. See the planning guide for your NonStop system to determine the types of disk drives it supports.

- “Internal SCSI Disk Drives”
- “M8xxx Fibre Channel Disk Drives”
- “Enterprise Storage System (ESS) Disks”

Internal SCSI Disk Drives

NOTE: Only NS16000 series systems support connections to internal SCSI disk drives installed in S-series I/O enclosures.

Internal SCSI disk drives are installed in NonStop S-series I/O enclosures.



CDT 106CDD

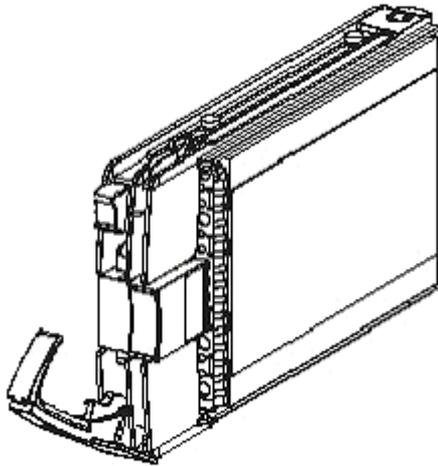
These disk drives are Class-1 CRUs.

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	See
Internal SCSI disk specifications	<i>NonStop S-Series Planning and Configuration Guide</i>
Internal SCSI disk commands	<i>SCF Reference Manual for the Storage Subsystem</i>
Classes of CRUs	<i>NonStop S-Series Planning and Configuration Guide</i>

M8xxx Fibre Channel Disk Drives

M8xxx Fibre Channel disk drives are installed in FCDM enclosures. A single disk drive is shown:



VST601.vxd

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	See
M8xxx Fibre Channel disk specifications	The planning guide for your NonStop system
M8xxx disk commands	<i>SCF Reference Manual for the Storage Subsystem</i>

Enterprise Storage System (ESS) Disks

The Enterprise Storage System (ESS) is any of several models of HP storage 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	See
ESS disk drives	HP XP storage documentation
ESS models supported on NonStop NS-series systems and NonStop BladeSystems	Your HP representative

Serial Attached SCSI (SAS) Disks and Solid State Drives

To determine whether your NonStop system supports CLIMs and SAS disks attached to CLIMs, see the planning guide for your NonStop system.

NonStop BladeSystems with two Storage CLIMs support SAS disk drives installed in SAS disk drive enclosures. Each SAS disk drive enclosure holds up to twenty-five 2.5" SAS disk drives, and up to four SAS disk drive enclosures can be connected to the two CLIMs, for a maximum of one hundred SAS disk drives. Each SAS disk drive enclosure is connected directly to both CLIMs.

Effective with the H06.24/J06.13 RVU, solid state drives (SSDs) in SAS enclosures connected to G6 CLIMs are supported. Effective with the H06.23/J06.12 RVU, SSDs, as well as HDDs in SAS enclosures connected to G6 CLIMs, can be partitioned. OSM provides support for SSDs and

displays partitioning information for HDDs and SSDs. See the *SCF Reference Manual for the Storage Subsystem* for more information.

For information about	See
SAS disk drive enclosures and SAS disk drives	<i>NonStop BladeSystem Planning Guide</i>
Configuring and managing CLIMs and Cluster I/O Protocol devices	<i>Cluster I/O Protocols (CIP) Configuration and Management Manual</i>
SAS disk drive commands	<i>SCF Reference Manual for the Storage Subsystem</i>

Monitoring Disk Drives

For information about	See
Tools for monitoring status, space use, configuration, and performance of disk drives	Appendix B (page 222)
Displaying EMS events generated by storage devices and subsystems	Chapter 5 (page 77)
Monitoring with OSM	“Monitoring Disk Drives With OSM” (page 143)
Monitoring with SCF	“Monitoring Drives With SCF ” (page 145)
Monitoring with HP SIM and its plug-ins	“Using HP SIM” (page 55)

Monitoring Disk Drives With OSM

For SAS disk drives, the physical SAS disk drive and the corresponding logical drive show as two separate objects in the OSM Service Connection tree pane. The physical disk drives are displayed under the SAS Disk Enclosures > SAS Disk Enclosure container objects, as shown in [Figure 8](#). The logical disks are displayed under the CLIM Attached Disks object, as shown in [Figure 9 \(page 144\)](#).

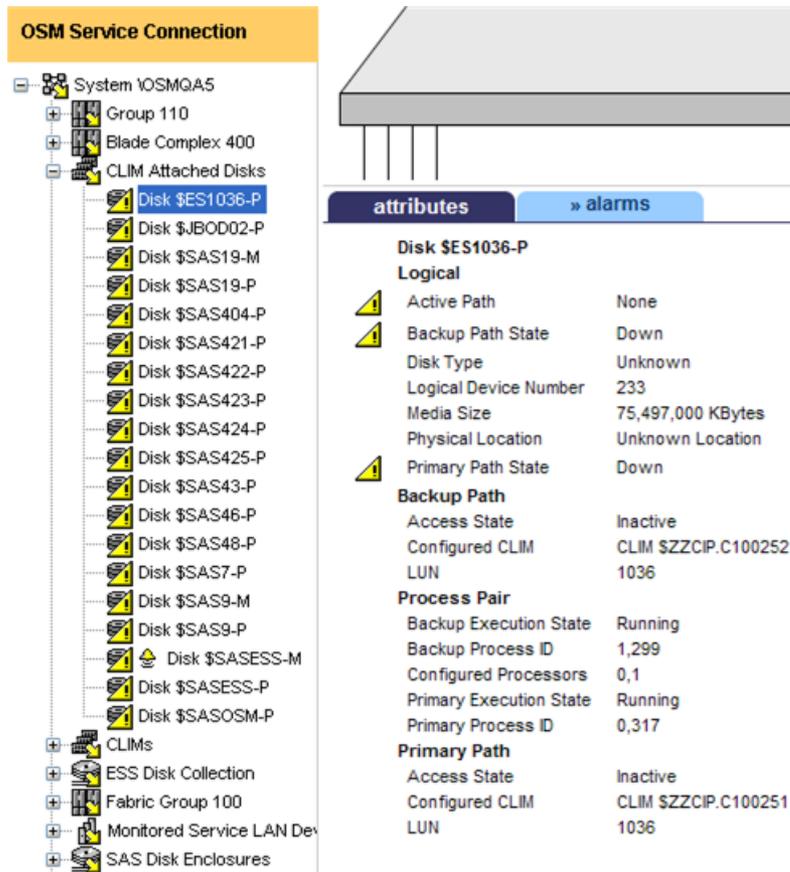
Figure 8 Physical Disk Drives in OSM

The screenshot displays the OSM Service Connection interface. On the left, a tree view shows the hierarchy: BladeCluster > System \BLOSM2 > Internal ServerNet X Fabric > Internal ServerNet Y Fabric > Group 110 > CLIM Attached Disks > CLIMs > Enclosure 100 > Monitored Service LAN Devices > SAS Disk Enclosures > SAS Disk Enclosure SGA736005B. The right pane shows the attributes for the selected enclosure:

SAS Disk Enclosure SGA736005B	
Physical	
Manufacturer	HP
Model	MSA70
Serial Number	SGA736005B
Connected CLIM1	
CLIM	SAS Card \$ZZCIP.C1002531.4
Logical	
Fan Status	Fully Redundant
Power Supply Status	Fully Redundant
Temperature Status	Normal
Physical Location	
Rack Name	BLOSM2
Rack Offset	38
Connected CLIM2	
CLIM	SAS Card \$ZZCIP.C1002533.4
Firmware	
Compare State	Same
Default File Name	\$\$SYSTEM.SYS00.fw0841
Default File Version	2.16
Version	2.16

VST201.vsd

Figure 9 Logical Disks in OSM



VST910.vsd

Task	See
Monitor the status of disk drives	<ul style="list-style-type: none"> • OSM Service Connection • OSM Event Viewer
Inventory the entire system, including disk drives	OSM Inventory View. You can save this view as an Excel file.
Use:	OSM Online Help
<ul style="list-style-type: none"> • OSM Service Connection • OSM Event Viewer • OSM Inventory View 	
Inventory multiple systems, including disk drives	OSM System Inventory Tool. You can save this inventory as an Excel file.
Use:	OSM Online Help (available within each application or tool)
<ul style="list-style-type: none"> • OSM Service Connection • OSM Event Viewer • OSM System Inventory Tool 	
For non-SAS disks, determine:	“Using the OSM Service Connection” (page 59)
<ul style="list-style-type: none"> • Service state • Primary path state • Backup path state 	

Task	See
<ul style="list-style-type: none"> Primary path access state Backup path access state <p>For SAS disks, check both the physical disk (under SAS Disk Enclosures) and the logical disk (under CLIM Attached Disks).</p> <p>For physical disk, determine the Device State.</p> <p>For logical disk, determine:</p> <ul style="list-style-type: none"> Service state Primary path state Backup path state Primary path access state Backup path access state <p>Check for alarms and degraded attribute values.</p> <p>For SAS disks, check both the physical disk (under SAS Disk Enclosures) and the logical disk (under CLIM Attached Disks) for alarms and degraded attribute values.</p>	<p>"Using the OSM Service Connection" (page 59)</p> <p>"Monitoring the State of Disk Drives" (page 148)</p>

Monitoring Drives With SCF

This section explains how to list disk volumes and determine their status.

- 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 Backup
        PID PID
        6 *STARTED STARTED *STARTED STARTED 0,257 1,257

STORAGE - Status DISK \COMM.$VIRCFG
LDev Primary Backup Mirror MirrorBackup Primary Backup
        PID PID
        146 *STARTED STARTED *STARTED STARTED 2,288 3,267

STORAGE - Status DISK \COMM.$WORK2
LDev Primary Backup Mirror MirrorBackup Primary Backup
        PID PID
        140 *STARTED STARTED *STARTED STARTED 5,278 4,273

STORAGE - Status DISK \COMM.$WEB02
LDev Primary Backup Mirror MirrorBackup Primary Backup
        PID PID
        143 *STARTED STARTED *STARTED STARTED 2,289 3,266

STORAGE - Status DISK \COMM.$ROOT
LDev Primary Backup Mirror MirrorBackup Primary Backup
        PID PID
        190 *STARTED STARTED *STARTED STARTED 3,268 2,287

STORAGE - Status DISK \COMM.$P1D02
LDev Primary Backup Mirror MirrorBackup Primary Backup
        PID PID
        247 *STARTED STARTED *STARTED STARTED 4,268 5,276

STORAGE - Status DISK \COMM.$P1D03
LDev Primary Backup Mirror MirrorBackup Primary Backup
        PID 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 PID	Backup 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"](#) (page 152).

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 PID	Backup 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 PID	Backup 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 PID	Backup PID
6	*STARTED	STARTED	*STARTED	STARTED	0,257	1,257

STORAGE - Status VIRTUAL DISK \COMM.\$VIEWPT

```

LDev   State      Primary Backup   Type   Subtype
      PID      PID
  147   STARTED    9,22   8,53   3     36

```

STORAGE - Status VIRTUAL DISK \COMM.\$WANA

```

LDev   State      Primary Backup   Type   Subtype
      PID      PID
  145   STARTED    8,77   9,56   3     36

```

STORAGE - Status VIRTUAL DISK \COMM.\$WEB

```

LDev   State      Primary Backup   Type   Subtype
      PID      PID
  144   STARTED    9,29   8,48   3     36

```

STORAGE - Status VIRTUAL DISK \COMM.\$WEBVPT

```

LDev   State      Primary Backup   Type   Subtype
      PID      PID
  142   STARTED    9,26   8,47   3     36

```

STORAGE - Status VIRTUAL DISK \COMM.\$WIPRO

```

LDev   State      Primary Backup   Type   Subtype
      PID      PID
  141   STARTED    9,27   8,51   3     36

```

STORAGE - Status VIRTUAL DISK \COMM.\$ZERO

```

LDev   State      Primary Backup   Type   Subtype
      PID      PID
  133   STARTED    8,78   9,57   3     36

```

STORAGE - Status VIRTUAL DISK \COMM.\$ZIMBU

```

LDev   State      Primary Backup   Type   Subtype
      PID      PID
  115   STARTED    9,28   8,52   3     36

```

STORAGE - Status DISK \COMM.\$VIRCFG

```

LDev   Primary Backup   Mirror MirrorBackup Primary Backup
      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 Backup
      PID          PID
  63 PRIMARY        ACTIVE   STARTED SubState 0,267 1,266
  63 BACKUP        INACTIVE STARTED SubState 0,267 1,266
  63 MIRROR        ACTIVE   STARTED SubState 0,267 1,266
  63 MIRROR-BACKUP INACTIVE STARTED SubState 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.TSYSDP2
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 27 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.

Table 27 Primary and Backup Path States for Disk Drives (continued)

Path State	Description
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 TAQL 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 (RWE): 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”](#) (page 153).

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 (RWEPR) : 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	
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 ” (page 151).

Table 28 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.
Defective tracks or sectors exist.	Output from SCF INFO DISK, BAD command indicates unspared defective sectors.
Disk errors exceed a certain limit.	Intm-errors-exceeded message
Slow I/O operations exceed a certain limit.	Slow-I/Os-threshold-exceeded message

SCSI Disk Drives Attached to NonStop NS-Series Systems

The most common disk drive problems on attached SCSI disk drives 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 NonStop NS-series system 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 29 Common Recovery Operations for Disk Drives

Problem	Recovery
Free-space fragmentation	Use the Disk Compression Program (DCOM) to consolidate disk space. See “Disk Compression Program (DCOM)” (page 222) .
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. For more information about these utilities, see “BACKUP” (page 222) and “Disk Space Analysis Program (DSAP)” (page 222) .
Down disk or disk path	“Recovery Operations for a Down Disk or Down Disk Path” (page 152) .

Table 29 Common Recovery Operations for Disk Drives *(continued)*

Problem	Recovery
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>. Disks come formatted from HP. No disk format utility is available. Return any disk that requires formatting to HP.</p>
Unspared defective sectors	<p>To check for unspared defective sectors with SCF:> INFO DISK \$*, BAD, SEL started, sub magneticTo check for unspared defective sectors with DSAP at a TACL prompt:> DSAP \$*</p> <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” (page 153).</p>
Performance problems	<p>Performance problems can have various causes, including path switches or a cache size that is too small. 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 Installation Guide for your NonStop system.</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_{nnn} and CSS_{nnn} 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 NonStop Support Center (GNSC) 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 Collection” (page 31) 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 Backup
              PID          PID
116 PRIMARY          ACTIVE     STARTED
116 BACKUP          INACTIVE  STARTED
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 Backup
              PID          PID
96  PRIMARY          ACTIVE     STARTED
96  BACKUP          INACTIVE  STARTED
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 Backup
              PID          PID
121 PRIMARY          INACTIVE  STOPPED    HARDDOWN   0,284 1,267
121 BACKUP          INACTIVE  STOPPED    HARDDOWN   0,284 1,267
121 MIRROR          ACTIVE     STARTED
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 (RWE): 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 Disk Drives

For information about	See
Complete syntax, examples, and considerations for the SCF commands used in this section	The planning guide for your NonStop system
Utilities such as: <ul style="list-style-type: none">• DCOM• DSAP• BACKUP• RESTORE	Guardian Disk and Tape Utilities Reference Manual
Other operations procedures involving disk drives	Guardian User's Guide
Using HP SIM and its plug-ins	<i>NonStop Cluster Essentials Installation and Quick Start Guide</i> <i>NonStop I/O Essentials Installation and Quick Start Guide</i> Online help for NonStop I/O Essentials and NonStop Cluster Essentials

13 Tape Drives: Monitoring and Recovery

- [“When to Use This Chapter” \(page 155\)](#)
- [“Overview of Tape Drives” \(page 155\)](#)
- [“Monitoring Tape Drives” \(page 155\)](#)
 - [“Monitoring Tape Drive Status With OSM” \(page 156\)](#)
 - [“Monitoring Tape Drive Status With SCF” \(page 157\)](#)
 - [“Monitoring Tape Drive Status With MEDIACOM” \(page 158\)](#)
 - [“Monitoring the Status of Labeled-Tape Operations” \(page 159\)](#)
- [“Identifying Tape Drive Problems” \(page 159\)](#)
- [“Recovery Operations for Tape Drives ” \(page 160\)](#)
 - [“Recovery Operations Using the OSM Service Connection” \(page 160\)](#)
 - [“Recovery Operations Using SCF” \(page 160\)](#)
- [“Related Reading for Tapes and Tape Drives” \(page 161\)](#)

When to Use This Chapter

This chapter provides an overview of operating, monitoring, and recovery operations for tape drives attached to NonStop BladeSystems or NonStop NS-series systems.

Overview of Tape Drives

For a listing of the tape drives supported by various NonStop systems and how the tape drives are attached, refer to the planning guide for your NonStop system.

Tape drives with a multimode fibre channel (MMF) 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 NonStop 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 a NonStop BladeSystem or NonStop NS16000 series system 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” \(page 155\)](#).

For further information, refer to the Planning Guide for your NonStop system.

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 10](#)) 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 10](#); for an IOMF2-connected tape drive, see the following figure.
 - 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 10 OSM: Monitoring Tape Drives Connected to an FCSA

The screenshot shows the OSM Service Connection interface. On the left is a tree view of system objects. The 'Tape Collection' object is expanded, showing several tape drives. 'Tape Drive \$LTO2' is selected and highlighted in blue. It has a yellow warning icon next to it. On the right, the 'attributes' tab is active, showing details for 'Tape Drive \$LTO2'. The details are organized into sections: 'Logical', 'Associated Device', and 'Controller Path'. The 'Device State' is 'Hard Down', 'Ready Status' is 'Not Ready', and 'Port Name' is 'Cx100000E0322312E1'. There is also an 'alarms' tab visible.

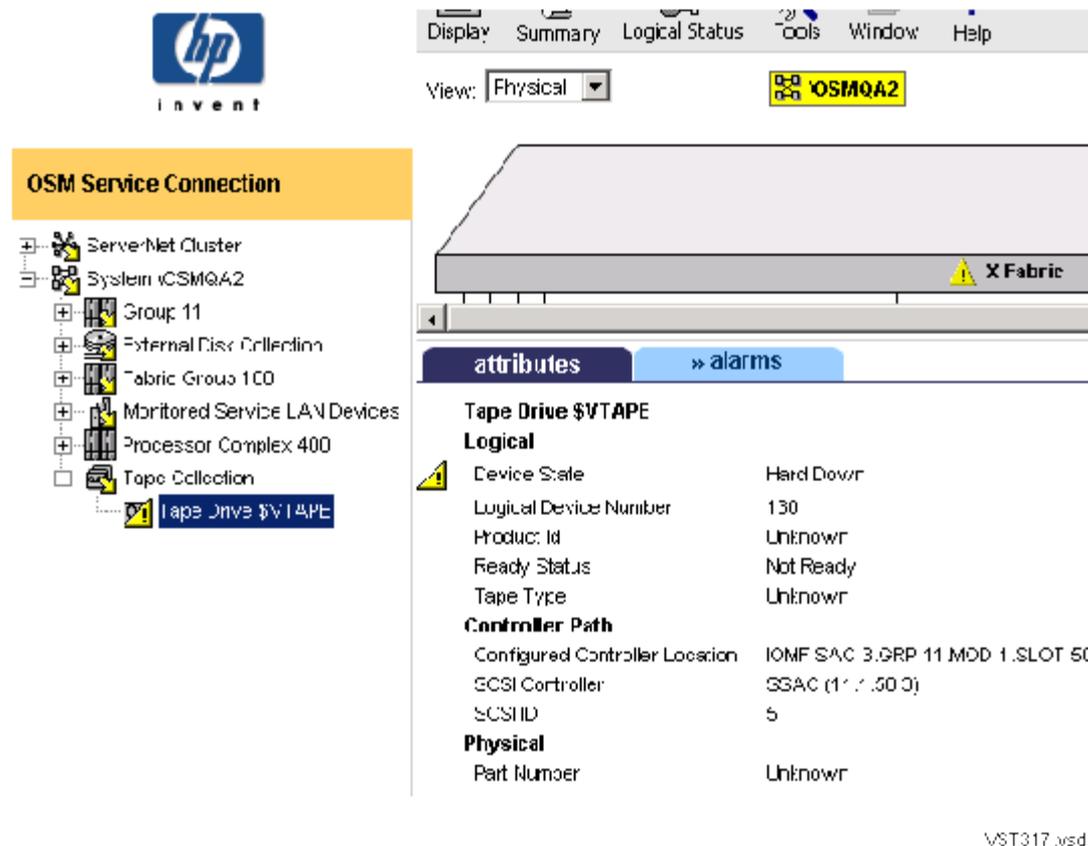
Tape Drive \$LTO2	
Logical	
Device State	Hard Down
Firmware Version	
Logical Device Number	199
Product Id	Unknown
Ready Status	Not Ready
Tape Type	Unknown
Associated Device	
Fibre Channel Router	FC Router 10
Controller Path	
Configured FSAC Location	FCSA.SAC-1.GRP:110.MOD:3.SLOT:2
FSAC	FCSA FSAC (110.0.2.1)
LUU	1
Port Name	Cx100000E0322312E1

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”](#) (page 160).
 - Service or replacement. Contact your service provider and refer to the [“Support and Service Collection”](#) (page 31) 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.

Figure 11 OSM: Monitoring Tape Drives Connected to an IOMF2



VST317.usd

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  Backup   DeviceStatus
      PID      PID
   93  STOPPED    1,287   0,279   NOT READY

STORAGE - Status TAPE \MINDEN.$TAPE0
LDev  State      Primary  Backup   DeviceStatus
      PID      PID
   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” (page 62) 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  Backup  DeviceStatus
      PID        PID
  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  Process Name
-----  Status  Name  Status  Type  Mode  -----
$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   Process
----- Drive  Status Name   Status Type   Mode   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 30 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 30 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 \$ZSVR, but it is not running. Tape operation is not allowed.	\$ZSVR 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 BACKUP 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 10 \(page 156\)](#)):
 - 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 10](#), 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 Tapes and Tape Drives

For more information about tapes and tape drives, refer to the documentation listed in [Table 31](#).

Table 31 Related Reading for Tapes and Tape Drives

For information about..	Refer to..
Tape drives	The planning guide for your NonStop system
BACKUP, RESTORE, and BACKCOPY utilities	<i>Guardian Disk and Tape Utilities Reference Manual</i> (for Enscribe and SQL/MP files)
BRCOM utility	<i>Backup and Restore 2.0 Manual</i> (OSS and SQL/MX files)
OSS pax utility	<i>Open System Services Management and Operations Guide</i> (backup and restore of OSS files)
Performing system operations involving tape drives	<i>Guardian User's Guide</i>
Replacing tape drives	"Support and Service Collection" (page 31)
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>
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>

14 Printers and Terminals: Monitoring and Recovery

- “When to Use This Chapter” (page 162)
- “Overview of Printers and Terminals ” (page 162)
- “Monitoring Printer and Collector Process Status” (page 162)
 - “Monitoring Printer Status” (page 162)
 - “Monitoring Collector Process Status” (page 163)
- “Recovery Operations for Printers and Terminals” (page 163)
 - “Recovery Operations for a Full Collector Process” (page 163)
- “Related Reading for Printers” (page 163)

When to Use This Chapter

This chapter 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 for Printers](#)” (page 163).

Overview of Printers and Terminals

Printers and terminals are attached to the NonStop system 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 section 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 for Printers](#)” (page 163).

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 Printers

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*

15 Applications: Monitoring and Recovery

- “When to Use This Chapter” (page 165)
- “Monitoring TMF” (page 165)
 - “Monitoring the Status of TMF” (page 165)
 - “Monitoring Data Volumes” (page 166)
 - “TMF States” (page 167)
- “Monitoring the Status of Pathway” (page 167)
 - “PATHMON States” (page 168)
- “Related Reading for Pathway” (page 169)

When to Use This Chapter

This chapter 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 section 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 TAACL 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”](#) (page 167).

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 32](#).

Table 32 TMF States

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: Services Waiting for Network Transactions to be Resolved Data Volumes Running Backout
Started	The TMF subsystem has started.
Stopped	The TMF subsystem is stopped.
Stopping	The TMF subsystem is stopping and is in one of these conditions: Waiting for Transactions to Finish Data Volumes Waiting for RDF 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 section 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  
SERVERCLASSES                   13        5        18        0        PENDING  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.

4. 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 Pathway

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*

16 Power Failures: Preparation and Recovery

- “When to Use This Chapter” (page 170)
- “System Response to Power Failures ” (page 170)
 - “NonStop Cabinets (Modular Cabinets) ” (page 171)
 - “External Devices ” (page 171)
 - “ESS Cabinets” (page 172)
 - “Networking CLIMs” (page 172)
 - “Storage CLIMs” (page 172)
 - “SAS Disk Drive Enclosures” (page 172)
 - “Air Conditioning” (page 172)
- “Preparing for Power Failure” (page 172)
 - “Set Ride-Through Time” (page 173)
 - “Monitor Power Supplies” (page 173)
 - “Monitor Batteries” (page 173)
 - “Maintain Batteries” (page 173)
- “Power Failure Recovery” (page 174)
 - “Procedure to Recover From a Power Failure” (page 174)
 - “Setting System Time” (page 174)
- “Related Reading for Power Failures” (page 174)

When to Use This Chapter

Use the chapter for information about how to prepare for power failures and how to recover if a power failure occurs. If you have a Telco system, contact your service provider for information about how to prepare for power failures and how to recover if a power failure occurs.

NOTE: For NonStop BladeSystems that shipped before J06.10, there are new recommended configuration settings that your service provider can implement to help avoid startup problems that might occur following a power outage. For more information, refer your service provider to *Configuring NonStop NS2200 Series and NS2100 Systems and NonStop BladeSystems to Avoid System Startup Problems*, located in the internal Service Procedures collection of NTL.

System Response to Power Failures

OSM Power Fail Support

When properly configured, OSM power fail support for NonStop NS-series systems and NonStop BladeSystems includes detection and monitoring of an outage and, if conditions dictate, a controlled shutdown of I/O operations and processors.

This shutdown does not include stopping or powering off the system, nor does it stop TMF or other applications. (HP encourages customers to execute scripts to shut down database activity before the processors are shut down.) The goal is completion of disk writes for I/Os that passed to disk

controllers before the ServerNet was shut down in order to achieve a relatively clean shutdown. This makes TMF recovery less time-consuming and difficult than if all power failed, the NonStop OS crashed, and disk writes did not complete.

NOTE: OSM power fail support works as described only after it has been properly configured, as described in [“Configure OSM Power Fail Support” \(page 173\)](#).

When OSM determines that all P-switches or ServerNet switches are running on UPS power, OSM recognizes this as a power outage and logs an event indicating the beginning of the ride-through time period, which was specified during configuration. OSM monitors whether AC power returns before the ride-through period ends, and acts accordingly:

- If AC power is restored before the ride-through period ends, the ride-through countdown is terminated, and OSM does not take further steps to prepare for an outage.
- If AC power is not restored before the ride-through period ends, OSM broadcasts a PFAIL_SHOUT message to processors (the processor running OSM is the last one in the queue) to shut down the system's ServerNet routers and processors in a fashion designed to allow completion of disk writes for items in transit through controllers and disks.

NOTE: Do not turn off the UPS as soon as the NonStop OS is down. The UPS continues to supply power until that supply is exhausted, and that time must be long enough for disk controllers and disks to complete disk writes.

NonStop Cabinets (Modular Cabinets)

NonStop NonStop systems 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 system 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 Planning Guide for your NonStop system.

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 NS16000 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.

Networking CLIMs

During a power failure, Networking CLIMs will not be put into a low-power mode. Shutdown of the CLIM may not be graceful. Refer to the *Cluster I/O Protocols (CIP) Configuration and Management Manual* for more information about how to shut off the CLIM.

Storage CLIMs

During a power failure, a Storage CLIM without Write Cache Enable will not be put into low-power mode. Shutdown of the CLIM may not be graceful. External modular disk and tape subsystems attached to the CLIM do not maintain power. 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.

A Storage CLIM with the Write Cache Enable option selected and a local UPS on that rack can preserve power long enough for the cached data to be flushed to disk. See the *SCF Reference Manual for the Storage Subsystem* for information on setting the WRITECACHE disk attribute to ENABLED (the default is DISABLED).

SAS Disk Drive Enclosures

SAS disk drive enclosures respond to a power failure as described in “[NonStop Cabinets \(Modular Cabinets\)](#)” (page 171). If a disk volume is configured with the WRITECACHE disk attribute enabled to allow for faster writes to SAS disks, HP recommends that the volume be connected to a UPS to give the drives enough time to flush the write cache in the event of power loss.

All disk drives in SAS enclosures connected to G6 CLIMs, including any partitions they might have, are supported by OSM. OSM will create an alarm and send a problem incident report when a SAS solid state drive (SSD) approaches or has reached its wear-out threshold.

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

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 Planning Guide for your NonStop system.

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 NonStop 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.

NOTE: For NonStop BladeSystems, the ride-through time must be actively set to at least 3 minutes to allow for possible loss of SNMP traps.

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

For NonStop BladeSystems, you must also configure OSM power fail support for Enclosure 100 (and only for Enclosure 100) by performing a **Configure Power Source as UPS** action on one power supply unit and performing a **Configure Power Source as AC** action on the other power supply unit in the enclosure. These actions are located under the Enclosure object. Configuring power fail support in this manner allows OSM, in the event of a power outage, to monitor the outage and, if the AC power is not restored before the specified ride-through time period expires, to initiate an orderly system shutdown.

For NonStop NS-series systems, 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 NonStop NS16000 series systems) or one IOAM or VIO power supply unit (for NonStop NS14000 series, NS1200, 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 Planning Guide for your NonStop system 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 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.

Power Failure Recovery

After a power failure, if AC power is restored to a NonStop system while the batteries are still holding up the system, it will not be necessary to restart the system.

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. Powering on a CLIM in a NonStop BladeSystem should automatically restart it and the CIP subsystem.

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 ” (page 176).
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 (or by a supercapacitor in the ServerNet switch in a c7000 enclosure in a NonStop BladeSystem) 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 Power Failures

For more information about preparing for and recovering from power failures:

- The effect of power failures on NonStop systems, see the Planning Guide for your NonStop system.
- 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.
- Write caching for SAS disks, see the *SCF Reference Manual for the Storage Subsystem*.

17 Starting and Stopping the System

- “When to Use This Chapter” (page 176)
- “Powering On a System ” (page 176)
 - “Powering On the System From a Low-Power State” (page 176)
 - “Powering On the System From a No Power State” (page 177)
 - “Power Cycling a NonStop BladeSystem Processor” (page 179)
 - “Power Cycling a NonStop NS2000 Series, NS2100, NS2200 Series, NS2300, or NS2400 Series Processor” (page 181)
- “Starting a System” (page 182)
 - “Loading the System” (page 183)
 - “Starting Other System Components” (page 185)
 - “Performing a System Load ” (page 186)
 - “Performing a System Load From a Specific Processor” (page 187)
 - “Reloading Processors” (page 188)
- “Minimizing the Frequency of Planned Outages” (page 189)
 - “Anticipating and Planning for Change” (page 190)
- “Stopping Applications, Devices, and Processes” (page 190)
- “Stopping the System” (page 191)
 - “Alerts” (page 191)
 - “Halting All Processors Using OSM” (page 191)
- “Powering Off a System” (page 192)
 - “System Shutdown Using OSM” (page 192)
 - “System Shutdown Using SCF” (page 193)
 - “Emergency Power-Off Procedure” (page 193)
- “Troubleshooting and Recovery Operations ” (page 193)
 - “Fans Are Not Turning” (page 193)
 - “System Does Not Appear to Be Powered On” (page 193)
 - “Green LED Is Not Lit After POSTs Finish” (page 194)
 - “Amber LED on a Component Remains Lit After the POST Finishes” (page 194)
 - “Components Fail When Testing the Power ” (page 194)
 - “Recovering From a System Load Failure ” (page 194)
 - “Getting a Corrupt System Configuration File Analyzed” (page 195)
 - “Recovering From a Reload Failure ” (page 195)

- [“Exiting the OSM Low-Level Link” \(page 196\)](#)
- [“Opening Startup Event Stream and Startup TACL Windows” \(page 196\)](#)
- [“Related Reading for Starting and Stopping a System” \(page 197\)](#)

When to Use This Chapter

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.

If you have a Telco system, contact your service provider for information about how to start the system (perform a system load) or stop or power off the system.

Powering On a System

Powering on a system delivers AC power to the system cabinets. Fans on the c7000 enclosure, processor switches, processor Blade Elements, IOAM or VIO enclosures, and FCDM 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 ” \(page 68\)](#).

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” \(page 176\)](#)
- [“Powering On the System From a No Power State” \(page 177\)](#)

Powering On the System From a Low-Power State

NOTE: NonStop BladeSystems and NonStop NS2200 series, NS2100, and NS2000 series systems do not have a low-power state. This procedure does not apply to NonStop BladeSystems and NonStop NS2200 series, NS2100, and NS2000 series systems.

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. For NonStop NS-series systems only, 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 NonStop NS14000 series, NS1200, 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, your service provider can refer to the NonStop Hardware Installation Manual for your NonStop system. To power on the system subsequently:

1. Before you power on any system enclosures, power on stand-alone system consoles (if any), the external system devices, and any other devices you want started when the system starts. External system devices include tape devices, SAS disk drive enclosures, 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 NonStop BladeSystem or NonStop NS-series system 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 all of the following that are present in your NonStop system: c7000 enclosure, NonStop server blades, processor switches, processor Blade Elements, CLIMs, IOAM or VIO enclosures, and FCDM 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 system. For more information about status LEDs, see [“Using the Status LEDs to Monitor the System ” \(page 68\)](#).
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:
 - a. Locate the circuit breaker that controls half the power cords.
 - b. Switch this breaker off.
 - c. Check that all components are still operating.

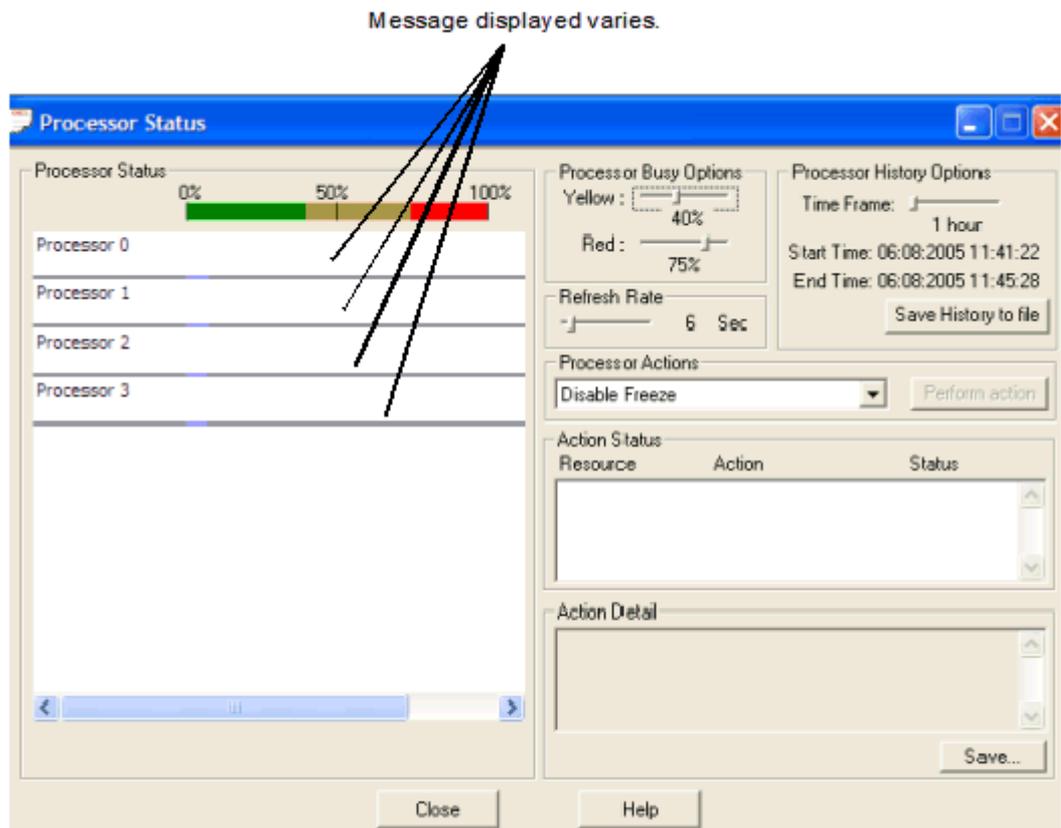
NOTE: The maintenance switch does not have redundant power.

- 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 ” \(page 194\)](#).
 - 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 in NonStop NS-series systems. For NonStop BladeSystems, logical processors down are displayed in Module 1 of their group, and there is no 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.

For NonStop BladeSystems and NonStop NS2400 series, NS2300, NS2200 series, NS2100, and NS2000 series systems, verify that all the processors have come up, as follows:

 - a. Close the current OSM Low-Level Link session, then launch and log on to a new session.
 - b. On the OSM Low-Level Link toolbar, click the **Processor Status** button to display the OSM Low-Level Link Processor Status dialog box.

Figure 12 Processor Status Dialog Box



- c. In the OSM Low-Level Link Processor Status dialog box, verify that all processors appear in the Processor Status list.
- d. Restart the OSM Low-Level Link application and then log on to the system.

NOTE: An already-running instance of the OSM Low-Level Link cannot be used here because it may show incorrect processor status. Instead, restart the OSM Low-Level Link application and log onto the system to start a new instance.

- e. Click **Processor Status**.
 - f. Verify that all processors are listed in the Processor Status dialog box.
 - If a NonStop BladeSystem processor is still not listed, power cycle it by using OA as described in [“Power Cycling a NonStop BladeSystem Processor”](#) (page 179).
 - If a NonStop NS2400 series, NS2300, NS2200 series, NS2100, or NS2000 series processor is still not listed, power cycle it as described in [“Power Cycling a NonStop NS2000 Series, NS2100, NS2200 Series, NS2300, or NS2400 Series Processor”](#) (page 181).
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”](#) (page 182).

Power Cycling a NonStop BladeSystem Processor

To use the Onboard Administrator (OA) to cycle power on a NonStop BladeSystem processor that has not come up, do the following:

1. Log into Onboard Administrator using the domain name of the OA of the HP NonStop BladeSystem c7000 enclosure that contains the processor that did not come up. The table below shows how the NonStop logical processors correspond to the slots in the enclosure.

NonStop Logical Processor	NonStop BladeSystem Enclosure	Slot in Enclosure
0	100	1
1	100	2
2	100	3
3	100	4
4	100	5
5	100	6
6	100	7
7	100	8
8	101	1
9	101	2
10	101	3
11	101	4
12	101	5
13	101	6
14	101	7
15	101	8

2. Click the Server Blade object that contains the processor that did not come up. The table above shows how the NonStop logical processors correspond to the slots in the enclosure.
3. Click the Virtual Devices tab.
4. On the Virtual Devices screen, click the **Press and Hold** virtual button to force the NonStop server blade power to turn off.
 - a. Wait 30 seconds to allow power to dissipate from the NonStop server blade before continuing with this procedure.
 - b. On the Virtual Devices screen, click the **Momentary Press** virtual button to restore power to the NonStop server blade.
 - c. Watch the power LED indicator on the Virtual Devices screen to verify that power was restored to the NonStop server blade:
 - If the green power light went on, go to Step 5.
 - If the green power light did not on, go to Step 4a to attempt to restore power again to the NonStop server blade.
 - If the green power light did not go on after three attempts to restore power to the NonStop server blade, contact your service provider.

5. Wait four minutes, and then use the OSM Low-Level Link application to verify the state of the affected processor, as follows:
 - a. Restart the OSM Low-Level Link application and then log on to the system.

NOTE: An already-running instance of the OSM Low-Level Link cannot be used here because it may show incorrect processor status. Instead, restart the OSM Low-Level Link application and log onto the system to start a new instance.

 - b. Click **Processor Status**.
 - c. Verify that the processor is listed in the Processor Status dialog box. If it is still not listed,
 - Go to Step 2 of this procedure and repeat Steps 2 through 5.
 - If it is still not listed after attempting Steps 2 through 5 a maximum of five times, power down the system, and then power it up again as described in [“Powering On the System From a No Power State”](#) (page 177).
6. When all processors are displayed in the OSM Low-Level Link, wait five minutes before starting your system as described in [“Starting a System”](#) (page 182).

NOTE: If you have performed this procedure, but are unable to have all installed processors displayed in the OSM Low-Level Link, contact your service provider.

Power Cycling a NonStop NS2000 Series, NS2100, NS2200 Series, NS2300, or NS2400 Series Processor

To use Integrated Lights-Out (iLO) to cycle power on a NonStop NS2400 series, NS2300, NS2200 series, NS2100, or NS2000 series processor that has not come up, follow these procedures.

- **For NonStop NS2000:**
 1. From the system console, use Internet Explorer to connect to the iLO IP address associated with the processor to be power cycled.
 2. Power cycle the processor as follows:
 - a. Click the **Virtual Devices** tab.
 - b. Click **Power and Reset** in the left pane.
 - c. Select **Power Off**.
 - d. Click **Submit**.
 - e. Click **OK** to dismiss the dialog box warning about shutting down the OS.
 - f. Wait 30 seconds.
 - g. Select **Power On**.
 - h. Click **Submit**.
 - i. Click **OK** to dismiss the dialog box asking if you wish to continue.
 3. Wait 4 minutes, and then use the OSM Low-Level Link application to verify the state of the affected processor:
 - a. Restart the OSM Low-Level Link application.
 - b. Click **Processor Status**.
 - c. Verify that the processor is listed in the Processor Status dialog box. If it is still not listed,
 - Go to Step 1 of this procedure and repeat the whole procedure.
 - If it is still not listed after attempting the procedure a maximum of five times, power down the system, and then power it up again as described in [“Powering On the System From a No Power State”](#) (page 177).
- **For NonStop NS2400 series, NS2300, NS2200 series, or NS2100:**

1. From the system console, use Internet Explorer to connect to the iLO IP address associated with the processor to be power cycled.
2. Power cycle the processor as follows:
 - a. Click **Power** and **Reset** in the left pane.
 - b. Click on **Momentary Press**.
 - c. Click **OK** to dismiss the dialog box asking if you wish to continue.
 - d. Wait until the Momentary Press button reappears for Power On.
 - e. Click **OK** to dismiss the dialog box asking if you wish to continue.
3. Wait 4 minutes, and then use the OSM Low-Level Link application to verify the state of the affected processor:
 - a. Restart the OSM Low-Level Link application.
 - b. Click **Processor Status**.
 - c. Verify that the processor is listed in the Processor Status dialog box. If it is still not listed,
 - Go to Step 1 of this procedure and repeat the whole procedure.
 - If it is still not listed after attempting the procedure a maximum of five times, power down the system, and then power it up again as described in [“Powering On the System From a No Power State” \(page 177\)](#).

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.

NonStop BladeSystems, NonStop N2400 series, NS2300, NS2200 series, and NS2100 systems must have access to DHCP, DNS, and BOOTP on the dedicated service LAN in order to get access the Halted State Services (HSS) firmware file on the dedicated service LAN needed to load the processors.

NS2000 and NS16000 systems with CLIMs attached require DHCP and DNS services as well, but not to access HSS files to boot; NS2000 and NS16000 systems only require DHCP and DNS because they use CLIM-attached disks.

DHCP, DNS, and BOOTP services and files can be located either on the two NonStop system consoles or on two designated CLIMs. The NonStop Maintenance LAN DHCP DNS Configuration Wizard is now used to configure and manage these services for your dedicated service LAN. Most of these configuration tasks are intended for service providers only, although there are tasks that any operator can perform. For more information, see the online help available within the wizard.

The NonStop Maintenance LAN DHCP DNS Configuration Wizard is installed as part of OSM Console Tools; upgrades are installed from the NonStop System Console Installer DVD. For the latest wizard enhancements and online help, it is important to get the most current version of the wizard available.

For more information on changes to the dedicated service LAN environment that can affect availability of these services and files, such as replacement of system consoles or CLIMs, your service provider can refer to:

- *Replacing a NonStop System Console* (located in the internal Service Procedures collection of NTL)
- *Replacing a CLIM* (launched by the Replace action in the OSM Service Connection; also located in the internal Service Procedures collection of NTL)
- *NonStop Dedicated Service LAN Installation and Configuration Guide* (located in the internal NonStop Hardware Service and Maintenance Publications collection of NTL)

Loading the System

For information about performing a system load to load the NonStop operating system, see the instructions for loading an RVU in the *Software Installation and Upgrade Guide* for your RVU.

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” \(page 183\)](#).

Normal System Load

Normally, you initiate the system load as described in [“Performing a System Load ” \(page 186\)](#). When you choose this method:

- Processor 0 or 1 is loaded. See [“System Load Paths for a Normal System Load” \(page 184\)](#).
- 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” \(page 203\)](#).

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 [Chapter 11 \(page 120\)](#).
 - 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” \(page 203\)](#).
- The Processor Element Dump Setting option becomes available in the System Load dialog box. For more information about this option, see [Chapter 11 \(page 120\)](#).
- You must initiate the load action as described in [“Performing a System Load From a Specific Processor” \(page 187\)](#).

System Load Disks

An NonStop 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) disk, SCSI disk, or CLIM-attached 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 FCDM enclosure connected to IOAM enclosure group 110:

	IOAM	FCSA	FCDM Enclosure			
Path	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 NonStop NS14000 series, NS1200, and NS1000 systems, Fibre Channel disks are connected to IOAMs or VIO enclosures located in group 100. For more information, see the Planning Guide for your NonStop NS14000 series, NS1200, or NS1000 system or the *Versatile I/O (VIO) Manual*.

- A load from \$SYSTEM attempts to load the system from a disk with default group, module, slot that varies with the system type. See the Planning Guide for your NonStop system to find the default startup characteristics for your system.
- A CLIM attached-load attempts to load the system from a disk in a SAS disk drive enclosure. CLIM-attached disks saved as System Load Configuration are displayed in the drop-down menu. CLIM attached-load is the only type of system load available for NonStop N2400 series, NS2300, NS2200 series, NS2100, and NS2000 series systems.
- You can configure additional alternate system disks to load from. To create an alternate system disk, see the Installation Guide for your NonStop system. Then, after using the Save as System Load Configuration action on the Disk object in the OSM Service Connection, this new configuration will be listed among the options under the Configuration Drop-down menu in the System Load dialog box of the OSM Low-Level Link.

System Load Paths for a Normal System Load

16 paths are available for loading. [Table 33 \(page 184\)](#) 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” \(page 193\)](#).

Table 33 System Load Paths in Order of Use

Data Travels				
Load Path	Description	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
6	Mirror	\$SYSTEM-M	0	Y
7	Mirror backup	\$SYSTEM-M	0	X

Table 33 System Load Paths in Order of Use *(continued)*

Data Travels				
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”](#) (page 194)

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. CLIMs should start automatically when power is applied.
- 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 for Starting and Stopping a System”](#) (page 197).

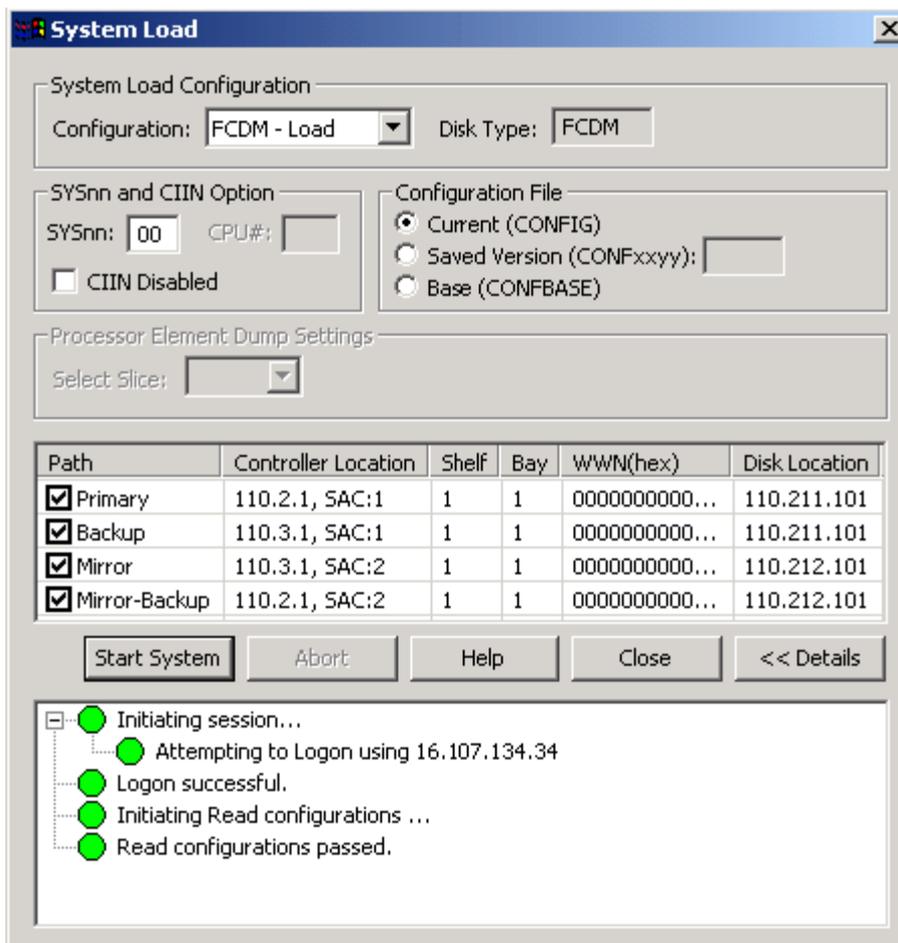
- You can include commands in startup command files that you invoke from a TACL prompt or another startup file. For some techniques to make startup command files run as efficiently as possible, refer to [Writing Efficient Startup and Shutdown Command Files](#) [“Writing Efficient Startup and Shutdown Command Files”](#) (page 206).

Performing a System Load

To perform a normal system load:

- Verify that all processors are halted as described in [“Stopping the System”](#) (page 191). All processors in the system must be halted before you initiate a system load.
- Log on to the OSM Low-Level Link.
- From the OSM Low-Level Link toolbar, click **Start system**.
- 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 13 System Load Dialog Box



- 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. In NonStop systems with suitable CLIMs, you can also select CLIM attached-Load to display a drop-down list of CLIM-attached SAS disks saved with System Load Configuration.
- 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.

3. In the Configuration File box, select a system configuration file. In most cases, you should select the Current (CONFIG) file.
4. 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.
5. To make changes to the load paths, double-click on a row in the Path window.
6. Click **Start system**.
7. For NonStop NS-series systems, only, 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)
8. After the System Load dialog box displays the “System Startup Complete” message, close the dialog box.
9. 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” \(page 188\)](#).

Performing a System Load From a Specific Processor

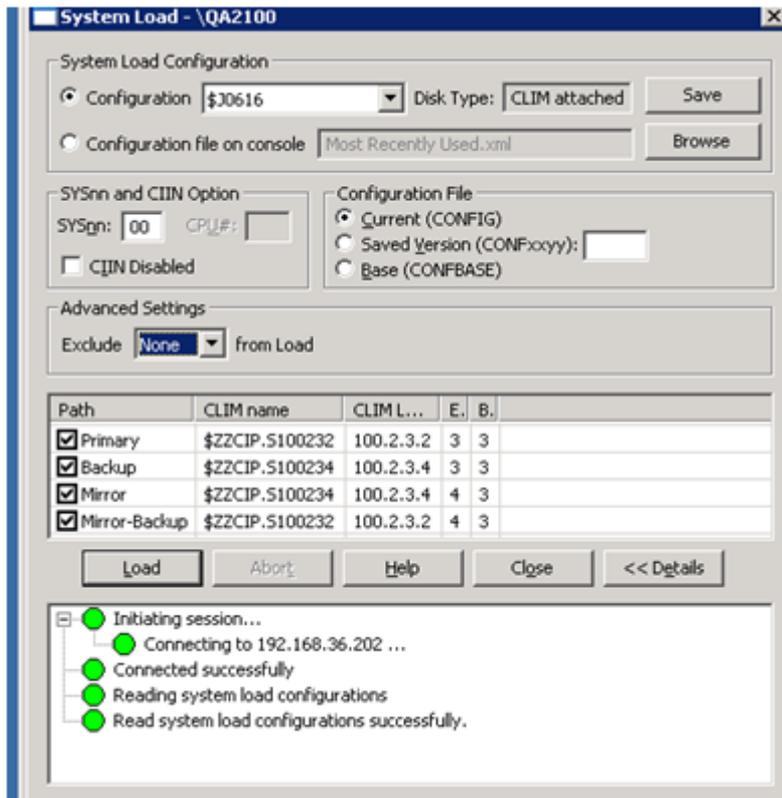
Use this method if you need to dump processor memory. See [Chapter 11 \(page 120\)](#). To perform a system load into a specified processor:

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, CLIM-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), SCSI disk, or CLIM. 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” \(page 185\)](#).
6. Select or clear the CIIN disabled check box. If you need to dump processors, you must select the CIIN disabled option.
7. For NonStop NS16000 series and NS14000 series systems, if you need to perform a memory dump of the logical processor you are loading, you can omit a Blade Element to prevent the processor element selected from being loaded or primed. Select the Blade Element from the **Exclude from load** drop-down menu in the Advanced Settings field of the System Load Configuration dialog box as shown in [Figure 14 \(page 188\)](#). See [“Dumping a Processor Using RCVDUMP for NonStop NS16000 Series and NS14000 Series Systems” \(page 136\)](#) for instructions on dumping the individual PE.

Figure 14 System Load Configuration Dialog Box



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 [Chapter 11 \(page 120\)](#).
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” (page 188)
- “Reloading Processors Using OSM” (page 189)

To reload a halted processor and perform memory dumps, use the reload procedures in [Chapter 11 \(page 120\)](#).

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*.

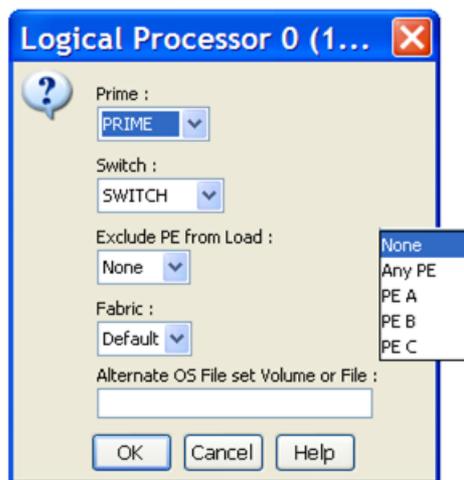
Reloading Processors Using OSM

The OSM Service Connection provides a Reload action on the Logical Processor object. You can perform the action on a single or multiple processors. For NonStop NS-series systems, 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 [Chapter 11 \(page 120\)](#).

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 \(page 189\)](#)), 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 Logical Processor Reload Parameters



VST014.vsd

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 Applications, 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” (page 191).

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 (such as a 6770 ServerNet cluster, 6780 ServerNet cluster, or BladeCluster), 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), the *ServerNet Cluster Manual* (for 6770 switches), or the *BladeCluster Solution Manual* (for BladeClusters).

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” \(page 191\)](#).
- 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” \(page 191\)](#).
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

You power off a NonStop system under special circumstances, such as before moving it to another physical location, prior to extensive service work, or simply to retire the system.

Use OSM or SCF to place most system components in a low-power state before you remove power to certain system components. Any attached I/O enclosures power off completely. The FCDM enclosures are not placed in a low-power state after the power off action or command. Some NonStop systems do not have low-power states.

There is no single programmatic command for powering off a NonStop system. You must perform these steps:

1. Stop applications. See [“Stopping Applications, Devices, and Processes” \(page 190\)](#).
2. Use either of these methods to place most system components in a low-power state:
 - [“System Shutdown Using OSM” \(page 192\)](#)
 - [“System Shutdown Using SCF” \(page 193\)](#)

NOTE: For NonStop BladeSystems, NS2400 series, NS2300, NS2200 series, NS2100, and NS2000 series systems, which do not have a low-power state, wait three minutes after performing Step 2 to allow cache to write to disks before you manually power off hardware components.

3. Locate the PDU circuit breakers. The number of PDUs depends on whether or not the rack has a c7000 enclosure. The number of circuit breakers on each PDU depends on the PDU model.
4. Switch all PDU circuit breakers to the OFF position. UPS output is connected to the PDU inputs. After switching the PDU circuit breakers OFF, the UPS is still on but its load is 0%.
5. Power down all UPSs and optional ERMs and disconnect each UPS from utility power. For accurate and detailed instructions about powering down a UPS, see the HP user guide for that UPS.
6. Locate the circuit breaker that controls the power cords and switch this breaker off.

System Shutdown 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**.
5. OSM displays this message:

This action will disable new I/O requests and halt processors. On an S-Series system, this action powers off S-Series enclosures. On an NS-Series system, this action will place some server components in low power mode. On a NonStop BladeSystem, none of the components will be powered off. After waiting for disk write cache to be written to media on SAS disk enclosures and ESS storage, you can power off all system components manually.

Do you still wish to execute this action?

6. Select Yes.

System Shutdown 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 Planning Guide for your NonStop system.

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 [Chapter 16 \(page 170\)](#).

Troubleshooting and Recovery Operations

Refer to the appropriate section for recovery information:

- [“Recovering From a System Load Failure ” \(page 194\)](#)
- [“Getting a Corrupt System Configuration File Analyzed” \(page 195\)](#)
- [“Recovering From a Reload Failure ” \(page 195\)](#)
- [“Exiting the OSM Low-Level Link” \(page 196\)](#)
- [“Opening Startup Event Stream and Startup TACL Windows” \(page 196\)](#)

If any of these problems occur when you power on a system, see the appropriate section for recovery information:

- [“Fans Are Not Turning” \(page 193\)](#)
- [“System Does Not Appear to Be Powered On” \(page 193\)](#)
- [“Green LED Is Not Lit After POSTs Finish” \(page 194\)](#)
- [“Amber LED on a Component Remains Lit After the POST Finishes” \(page 194\)](#)
- [“Components Fail When Testing the Power ” \(page 194\)](#)

Fans Are Not Turning

If the fans do not start turning a few seconds after you power on the system, 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” \(page 192\)](#). Contact your service provider.

System Does Not Appear to Be Powered On

If AC power is being supplied to the system and the fans are turning but the system still does not appear to be powered, the system might be running internal tests. Wait several minutes (at least

10 minutes for large configurations). If the system 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

NOTE: For more information on system startup issues specific to NonStop BladeSystems, see *Troubleshooting Tips for Powering On NonStop NS2100 Systems, NS2200 Series Systems and NonStop BladeSystems*, located in the Troubleshooting Tips collection of NTL.

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” \(page 185\)](#).
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” \(page 182\)](#). In the Configuration File box, select **Base (CONBASE)** as the configuration file.
 - b. Reload the remaining processors. See [“Reloading Processors” \(page 188\)](#).
 - c. From the Startup TACL window, configure a tape drive.
 - d. Restore a previously backed-up configuration file.
 - e. Load the system as described in [“Starting a System” \(page 182\)](#) 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 TACL 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” \(page 185\)](#).
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” \(page 134\)](#).
5. Try a soft reset of the processor.
6. Reload the processor or processors as described in [Chapter 11 \(page 120\)](#). 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

NOTE: OSM Low-Level Link version T0633AAT and later launches comForte MR-Win6530 for startup TACL and event stream windows. MR-Win6530 launches a single session per system load.

For information on installing and getting started with MR-Win6530, see the *NonStop System Console Installer Guide* for version S7X-SWV2/HNSC-SWV2 (H06.11) or later. To launch a TACL session directly from MR-Win6530, see the *MR-Win6530 User's Guide* (shipped with the console-based product).

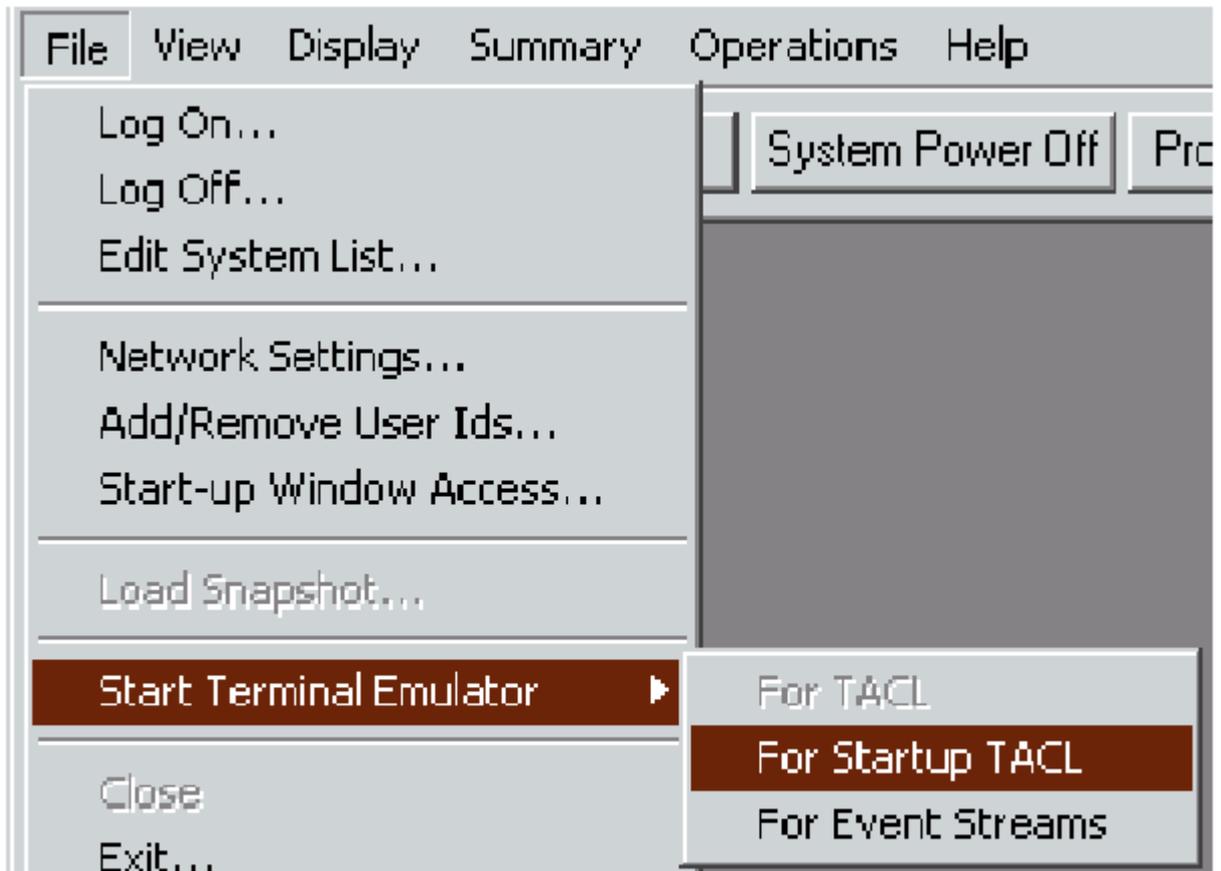
When you perform a normal system load, the Startup Event Stream and Startup TACL windows open automatically. If the windows do not open or if you close them, you can reopen them using either OSM or comForte MR-Win6530.

To open startup event stream windows and startup TACL windows from OSM:

1. Log on to the OSM Low-Level Link.

- From the File menu, select **Start Terminal Emulator > For Startup TACL**.

Figure 16 Opening a Startup TACL Window



VST813.vsd

To open startup event stream windows and startup TACL windows using comForte MR-Win6530:

- Select **Start > All Programs > MR-Win6530 > MR-Win6530**.
- Refer to the Win6530 Help or to the *Win6530 User Guide*, both available from **Start > All Programs > MR-Win6530 > MR-Win6530**.

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 Starting and Stopping a System

For more information, refer to the documentation listed in [Table 34](#).

Table 34 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 and J 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	Installation Guide for your NonStop system

Table 34 Related Reading for Starting and Stopping a System *(continued)*

For Information About	Refer to
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	Chapter 18 (page 199) and <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>
Starting and stopping CLIMs and the CIP subsystem	<i>Cluster I/O Protocols (CIP) Configuration and Management Manual</i>
Using the Onboard Administrator (OA)	<i>HP BladeSystem Onboard Administrator User Guide</i>

18 Creating Startup and Shutdown Files

This section describes command files that automatically start and shut down an NonStop system.

- “Automating System Startup and Shutdown” (page 200)
 - “NonStop Cluster Boot Application” (page 200)
 - “Managed Configuration Services (MCS)” (page 200)
 - “Startup” (page 201)
 - “Shutdown” (page 201)
 - “For More Information ” (page 201)
- “Processes That Represent the System Console” (page 201)
 - “\$YMIO.#CLCI” (page 201)
 - “\$YMIO.#CNSL” (page 201)
 - “\$ZHOME” (page 202)
 - “\$ZHOME Alternative” (page 202)
- “Example Command Files” (page 202)
- “CIIN File” (page 203)
 - “Establishing a CIIN File” (page 203)
 - “Modifying a CIIN File” (page 204)
 - “If a CIIN File Is Not Specified or Enabled in OSM” (page 204)
 - “Example CIIN Files” (page 204)
- “Writing Efficient Startup and Shutdown Command Files” (page 206)
 - “Command File Syntax” (page 206)
 - “Avoid Manual Intervention” (page 206)
 - “Use Parallel Processing” (page 206)
 - “Investigate Product-Specific Techniques” (page 207)
- “How Process Persistence Affects Configuration and Startup” (page 207)
- “Tips for Startup Files” (page 207)
- “Startup File Examples” (page 208)
 - “System Startup File” (page 208)
 - “Spooler Warm-Start File ” (page 209)
 - “TMF Warm-Start File” (page 210)
 - “TCP/IP Stack Configuration and Startup File” (page 210)
 - “CP6100 Lines Startup File” (page 212)
 - “X.25 Lines Startup File” (page 213)

- “Expand-Over-IP Line Startup File” (page 213)
- “Expand Direct-Connect Line Startup File” (page 214)
- “Tips for Shutdown Files” (page 214)
- “Shutdown File Examples” (page 214)
 - “System Shutdown File” (page 214)
 - “CP6100 Lines Shutdown File” (page 215)
 - “ATP6100 Lines Shutdown File” (page 215)
 - “X.25 Lines Shutdown File” (page 216)
 - “Printer Line Shutdown File” (page 216)
 - “Expand-Over-IP Line Shutdown File” (page 216)
 - “Direct-Connect Line Shutdown File” (page 216)
 - “Spooler Shutdown File” (page 217)
 - “TMF Shutdown File” (page 217)

Automating System Startup and Shutdown

NonStop Cluster Boot Application

HP NonStop Cluster Essentials includes a boot application that provides the functionality to boot and provision both NonStop and Linux systems together in sequence or in parallel, using a predefined boot configuration. Users can define one or many boot configurations and use them for subsequent boots by one click. The boot configurations can also be exported to other CMSs.

The steps in a boot configuration can be performed in parallel or in sequence. A step can be either running a script or performing the boot (system load) or delay. Users can specify whether the failure of a step should abort the boot operation or not. Any specified script can reside on the system or on the console. The console script can be run on the console itself, or transferred to the system and run there. For a NonStop system, a script can be either a Guardian-based script or an OSS-based script.

The boot application allows the user to orchestrate any simple or complex workflow for single or multiple NonStop or Linux systems.

HP NonStop Cluster Essentials has support for booting NonStop systems using Cluster I/O Module (CLIM) attached disks, Fibre Channel Disk Module (FCDM) disks, or SCSI disks. Users can also define global NonStop boot parameters (such as `SYSDISK`, system disk locations, and so on) and use them for individual systems.

For detailed instructions about using the NonStop Cluster Boot Application, see the *NonStop Cluster Boot Application User's Guide*, which is available in the online help for Nonstop Cluster Essentials.

Managed Configuration Services (MCS)

NonStop systems 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 Planning Guide for your NonStop system.

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” (page 203) 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 ..	Refer to..
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> and the <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 NonStop systems, 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 is configured with its INFILE, OUTFILE, and HOMETERM set to \$YMIOP.#CLCI. Messages written to \$ZHOME will appear on the \$YMIOP.#CLCI window if the Startup TACL window “Opening Startup Event Stream and Startup TACL Windows” (page 196) is open and paused. When the \$YMIOP.#CLCI window is closed, connectivity is lost and \$ZHOME will then write the message to \$0. Specifically, the \$ZHOME process writes the messages to its configured OUTFILE, which by default is \$YMIOP.#CLCI. If it gets an error writing to the OUTFILE, it switches to the collector \$0, which is hard-coded.

\$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 system:

- 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 Low-Level Link online help.

- 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 (RWEP) 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 `SYSnnn` subvolume and name it CIIN.

Modifying a CIIN File

After the CIIN file is established on `$SYSTEM.SYSnnn` (as part of running DSM/SCM), you can modify the contents of `SYSnnn.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
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 <code>\$YMIOP.#CLCI</code> started by the CIIN file) to complete the system startup process.
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.
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.
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.
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.

NOTE: 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.

NOTE: 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 ($\$S_n$) 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 Planning Guide for your NonStop system.

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 [Chapter 17 \(page 175\)](#). For information about automating disk processes upon startup, see the Planning Guide for your NonStop 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 more information, see [“Example Command Files” \(page 202\)](#).

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.STRTLP, 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 %%hFFFFFF00
+ 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.STRATAP, OUT $ZHOME /

== This is $SYSTEM.STARTUP.STRATAP

== 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” \(page 202\)](#).

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
```

19 Preventive Maintenance

- “When to Use This Chapter” (page 218)
- “Monitoring Physical Facilities” (page 218)
 - “Checking Air Temperature and Humidity” (page 218)
 - “Checking Physical Security” (page 218)
 - “Maintaining Order and Cleanliness” (page 218)
 - “Checking Fire-Protection Systems” (page 219)
- “Cleaning System Components” (page 219)
 - “Cleaning an Enclosure” (page 219)
 - “Cleaning and Maintaining Printers” (page 219)
 - “Cleaning Tape Drives” (page 219)
- “Handling and Storing Cartridge Tapes” (page 220)

When to Use This Chapter

This chapter describes routine maintenance tasks required for NonStop systems.

Monitoring Physical Facilities

This section 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 section 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.

NOTE: 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.

NOTE: 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, H-Series, and J-Series RVUs

Users familiar with systems running G-series RVUs will find several major differences in the operational environment of systems running H-series and J-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 and J-series RVUs, these changes have been made:

- TSM does not support H-series or J-series; you must use OSM as the system management tool for NonStop NS-series systems and NonStop BladeSystems.
- In power failures, there is no memory hold-up for H-series and J-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 and J-series.
- In H-series and J-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 and J-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 and J-series systems, `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 and J-series, DSM/SCM installation default is Manage OSS Files. For G-series, the default is not to manage OSS files.
- In H-series and J-series, 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 and J-series because of the new debuggers and debugger commands.
- The H-series and J-series OSS environment does not support TNS execution. OSS programs must be migrated to TNS/E native mode to run on an H-series or J-series system.

B Tools and Utilities for Operations

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 NonStop system. 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, see [“Related Reading for Tools and Utilities”](#) (page 227).

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.

CLIMCMD

The CLIMCMD mechanism executes CLIM commands from the NonStop host system TA CL.

CLIMCONFIG

The climconfig command is used to configure network resources on a CLIM. It works in much the same way that SCF SUBNET configuration works in other NonStop TCP/IP products.

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 fragments. 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 TA CL 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 NonStop system. You can use FUP to create, display, and duplicate files; load data into files; alter file characteristics; and purge files.

HP Systems Insight Manager (HP SIM)

HP Systems Insight Manager (HP SIM) provides a very high-level view of overall system health for multiple NonStop systems and other systems in a single window. To get greater detail about most issues and to service the NonStop system, use OSM.

Insight Control Power Management

Insight Control Power Management is an HP SIM plug-in application that provides centralized monitoring and control of server power consumption and thermal output, sets Power Regulator policies for groups of servers, and optimizes and controls usage of data center power and cooling facilities.

Use Insight Control Power Management to set the Power Regulator mode in NonStop BladeSystem NB54000c, NB54000c-cg, NB56000c, or NB56000c-cg after Power Regulator is enabled on the System object in the OSM Service Connection.

Insight Remote Support Advanced

HP Insight Remote Support Advanced replaces the OSM Notification director in both modem-based and HP Instant Support Enterprise Edition (ISEE) remote support solutions. For more information about Insight Remote Support Advanced, refer to the HP Integrity NonStop Service information collection of NTL.

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 Cluster Essentials

NonStop Cluster Essentials can define a cluster of NonStop (and even Linux) systems so that you can perform actions on multiple systems with a single action. Combine NonStop Cluster Essentials with NonStop I/O Essentials to use a single action to perform actions on CLIMs in different systems.

NonStop Cluster Performance Essentials

NonStop Cluster Performance Essentials is a plug-in to NonStop Cluster Essentials that provides integrated performance monitoring of homogeneous clusters of NonStop systems or heterogeneous clusters of NonStop and Linux systems.

NonStop I/O Essentials

The NonStop I/O Essentials product is an HP SIM plug-in that virtualizes the configuration and control of all CIP (Cluster I/O Protocols) tasks by providing a graphical user interface alternative to the command line interfaces of the CLIMCMD tool and SCF. Combine NonStop I/O Essentials with NonStop Cluster Essentials 2.0 (or later) to use a single action to perform actions on CLIMs in different systems.

NonStop Maintenance LAN DHCP DNS Configuration Wizard

The NonStop Maintenance LAN DHCP DNS Configuration Wizard can be used to configure DHCP, DNS, and BOOTP servers. DHCP and DNS servers are required for NonStop systems and the CLIMs connected to these systems. BOOTP servers are only for the J-series NonStop BladeSystems, NonStop NS2400 series, NS2300, NS2200 series, and NS2100 systems and make the Halted State Services (HSS) files available for processors to boot from.

The NonStop Maintenance LAN DHCP DNS Configuration Wizard replaces the CLIM Boot Service Configuration Wizard. It is delivered and installed on the system console as part of the OSM Console Tools product, version T0634 ABB and later. For information on pre-requisites and installing the wizard, see the *NonStop System Console Installer Guide*.

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.

NonStop Software Essentials

NonStop Software Essentials is an HP SIM plug-in that provides software installation and management functionality for NonStop systems. NonStop Software Essentials uses DSM/SCM server software and provides a replacement for the DSM/SCM Planner Interface tasks as well as certain Host Maintenance Interface functions that is more secure, is easier to use, and overcomes other shortcomings of DSM/SCM. NonStop Software Essentials coexists with DSM/SCM; you can use either NonStop Software Essentials or DSM/SCM for tasks in which there is overlapping functionality. For more information, see the *NonStop Software Essentials Installation and Quick Start Guide* and the online help for NonStop Software Essentials.

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.

Onboard Administrator (OA)

The HP BladeSystem Onboard Administrator (OA) provides details about alarms on NonStop BladeSystem c7000 enclosure components. Invoke the HP BladeSystem Onboard Administrator (OA) URL via an action on the enclosure object in the OSM tree view.

OSM Product Suite

The HP Open System Management (OSM) product is the required system management tool for NonStop NS-series systems and NonStop BladeSystems; they are not supported by TSM. For NonStop BladeSystems, OSM works together with the Onboard Administrator (OA) and Integrated Lights-Out (iLO) management interfaces to manage c7000 enclosures. OSM also works with HP SIM to allow HP SIM to manage NonStop systems alongside other HP servers and storage from a single user interface.

For more information on the OSM package, including a description of the individual applications, see the *OSM 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.

HP Tandem Failure Data System (TFDS)

HP Tandem Failure Data System (TFDS) automates many tasks associated with data collection and resource recovery in the event of software, processor, and subsystem failures. In response to a failure, TFDS collects the information needed for problem identification and resolution during subsequent analysis and repair efforts. When a failure results in a processor halt, TFDS automatically reloads and dumps the halted processor to disk.

TFDSCOM

TFDSCOM is the user interface for TFDS. You can make changes to the TFDS configuration by using the TFDSCOM commands, which are described in the *HP Tandem Failure Data System (TFDS) Manual*.

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 system 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.

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 Tools and Utilities

For more information about tools and utilities used for system operations, refer to the documentation listed in [Table 35](#).

Table 35 Related Reading for Tools and Utilities

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, H-series, and J-series RVUs.
Down System CLIM Firmware Update Tool	Down System CLIM Firmware Update Tool online help	This online help system describes the Down System CLIM Firmware Update Tool, which is delivered and installed on the system console as part of the OSM Console Tools product.
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.
HP Systems Insight Manager (HP SIM)	Online help <i>HP SIM for NonStop Manageability</i>	This online help system describes how to view overall system health for multiple NonStop systems and other systems in a single window. This manual describes how to use HP SIM to manage NonStop systems.
Insight Control Power Management	<i>HP SIM for NonStop Manageability</i> <i>Insight Control Power Management User Guide</i>	This manual describes how to use HP SIM to manage NonStop systems. It includes using Insight Control Power Management to set Power Regulator modes. This manual describes Insight Control Power Management in general, but is not specific for NonStop systems.
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>	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.

Table 35 Related Reading for Tools and Utilities *(continued)*

Tool	Documentation	Description
	<i>DSM/Tape Catalog Operator Interface (MEDIACOM) Manual</i>	This manual explains how to run a MEDIACOM session and describes the purpose and the syntax of the MEDIACOM commands.
	<i>Guardian User's Guide</i>	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.
NonStop Cluster Essentials	<i>NonStop Cluster Essentials Installation and Quick Start Guide</i>	Describes how to install, configure, and get started with the NonStop Cluster Essentials product.
	NonStop Cluster Essentials online help	Describes how use the NonStop Cluster Essentials product.
NonStop I/O Essentials	<i>NonStop I/O Essentials Installation and Quick Start Guide</i>	Describes how to install, configure, and get started with the NonStop I/O Essentials product.
	NonStop I/O Essentials online help	Describes how use the NonStop I/O Essentials product.
NonStop Maintenance LAN DHCP DNS Configuration Wizard	NonStop Maintenance LAN DHCP DNS Configuration Wizard online help	This online help system describes the NonStop Maintenance LAN DHCP DNS Configuration Wizard, which is delivered and installed on the system console as part of the OSM Console Tools product
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.
NonStop Software Essentials	<i>NonStop Software Essentials Installation and Quick Start Guide</i>	Describes how to install, configure, and get started with the NonStop Software Essentials product.
	NonStop Software Essentials online help	Describes how use the NonStop Software Essentials product.
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.
Onboard Administrator (OA)	<i>HP BladeSystem Onboard Administrator User Guide</i>	This guide describes how to use the Onboard Administrator to manage NonStop BladeSystems.
OSM package	<i>OSM Service Connection User's Guide</i> (available in NTL and as online help within OSM Service Connection)	This guide describes how to use the OSM Service Connection to monitor and perform actions on system and cluster resources.
	<i>OSM Configuration Guide</i>	This guide includes information on OSM: <ul style="list-style-type: none"> • Server processes • Configuration options • Applications and tools
	<i>NonStop System Console Installer Guide</i>	This guide describes how to install OSM client-based components and other required system console software.

Table 35 Related Reading for Tools and Utilities *(continued)*

Tool	Documentation	Description
Online help	<p>Online help is also available from within each of these OSM applications:</p> <ul style="list-style-type: none"> • NonStop Maintenance LAN DHCP DNS Configuration Wizard • Down System CLIM Firmware Update Tool • OSM Low-Level Link • OSM Notification Director • OSM Event Viewer • OSM System Inventory Tool • OSM Certificate Tool • OSM Emulator File Converter • OSM guided procedures (within the OSM Service Connection) 	
PATHCOM	<i>TS/MP System Management Manual</i>	<p>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.</p>
PEEK	<i>PEEK Reference Manual</i>	<p>This manual describes PEEK, a utility used to monitor statistical data about processors.</p>
RESTORE	<i>Guardian Disk and Tape Utilities Reference Manual</i>	<p>This manual describes these disk and tape utilities: BACKCOPY, BACKUP, DCOM, DSAP, and RESTORE. This manual supports D-series, G-series, H-series, and J-series RVUs.</p>
SPOOLCOM	<p><i>Spooler Plus Utilities Reference Manual</i></p> <p><i>Guardian User's Guide</i></p>	<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	Planning Guide for your NonStop system	<p>This guide describes how to automate startup and shutdown procedures.</p>

Table 35 Related Reading for Tools and Utilities *(continued)*

Tool	Documentation	Description
Subsystem Control Facility (SCF)	<i>SCF Reference Manual for H and J Series RVUs</i>	This manual describes the operation of SCF on H and J 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 a NonStop system. 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.
TFDS	<i>HP Tandem Failure Data System (TFDS) Manual</i>	This manual describes TFDS and TFDS.COM.
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.
Web ViewPoint	<i>Web ViewPoint Version 5.2 User Guide</i>	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: <ul style="list-style-type: none"> • Monitors and displays EMS events • Identifies and lists all supported subsystems • Manages NonStop system subsystems and user applications in a secure, automated, and customizable way

Table 35 Related Reading for Tools and Utilities *(continued)*

Tool	Documentation	Description
ViewPoint	<i>ViewPoint Manual</i>	<ul style="list-style-type: none"><li data-bbox="1082 191 1511 254">• Monitors and graphs performance attributes and trends<li data-bbox="1082 260 1511 323">• Investigates and displays most active system processes<li data-bbox="1082 329 1511 392">• Offers simple navigation and a point-and-click command interface <p data-bbox="1082 401 1511 764">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. Although the ViewPoint Manual applies to D-series, G-series, H-series, and J-series RVUs, some information might apply only to D- and G-series RVUs.</p>
ViewSys	<i>ViewSys User's Guide</i>	<p data-bbox="1082 779 1511 869">This guide describes the operation of ViewSys and interpretation of the program output.</p>

D Converting Numbers

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 36](#) describes the binary, octal, decimal, and hexadecimal number systems.

Table 36 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 system, 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 OX 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 17 \(page 233\)](#).

Figure 17 Binary to Decimal Conversion

Placeholder values	...	64	32	16	8	4	2	1
Binary number	...			1	1	0	1	1

1 * 1	=	1
1 * 2	=	2
0 * 4	=	0
1 * 8	=	8
1 * 16	=	16
		27

CDT 607CDD

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 18 \(page 234\)](#).

Figure 18 Octal to Decimal Conversion

Placeholder values	...	4096	512	64	8	1	
Octal number	...		1	3	7	5	

$1 * 512$	$= 512$
$3 * 64$	$= 192$
$7 * 8$	$= 56$
$5 * 1$	$= 5$
	<hr/>
	765

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 19 \(page 235\)](#).

Figure 19 Hexadecimal to Decimal Conversion

Placeholder values	...	4096	256	16	1
Hexadecimal number	...	B	A	1	0

				$0 * 1$	= 0
				$1 * 16$	= 16
		$10 * 256$			= 2560
	$11 * 4096$				= 45056
					47632

C0T 609C00

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 KCC 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.

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. In a domestic environment this product may cause radio interference 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 (2004/108/EC) and the Low Voltage Directive (2006/95/EC) 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

Republic of Turkey Notice

Türkiye Cumhuriyeti: EEE Yönetmeliğine Uygundur.

Ukraine Addendum to User Documentation

Україна Додаток до документації користувача: Обладнання відповідає вимогам Технічного регламенту щодо обмеження використання деяких небезпечних речовин в електричному та електронному обладнанні, затвердженого постановою Кабінету Міністрів України від 3 грудня 2008 № 1057

Збережіть цей документ разом із документацією користувача для цього виробу

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 or Class 1M 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:

- LASER RADIATION - DO NOT VIEW DIRECTLY WITH OPTICAL INSTRUMENTS. CLASS 1M LASER PRODUCT.
- 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 HP NonStop Technical Library at <http://www.hp.com/go/nonstop-docs>. Under HP NonStop Technical Library, select **HP Integrity NonStop Server Safety and Compliance > Safety and Regulatory Information>Waste Electrical and Electronic Equipment (WEEE)**.

Important Safety Information

Safety information is available from the HP NonStop Technical Library at <http://www.hp.com/go/nonstop-docs>. Under HP NonStop Technical Library, select **HP Integrity NonStop Server Safety and Compliance > Safety and Regulatory Information>Safety and Compliance**. 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.

Table of Toxic and Hazardous Substances/Elements and Their Contents

(as required by China's Management Methods for Controlling Pollution by Electronic Products)

有毒、有害物质/元素及其含量表

依照中国《电子信息产品污染控制管理办法》

部件名称	有毒、有害物质和元素					
	铅 (Pb)	汞 (Hg)	镉 (Cd)	六价铬 (Cr(VI))	多溴联苯 (PBB)	多溴代二苯醚 (PBDE)
机箱/其它	X	0	0	0	0	0
印刷电路组件 (PCAs)	X	0	0	0	0	0
存储设备	X	0	X	0	0	0
电缆和配线	X	0	0	0	0	0
平板显示器	X	X	0	0	0	0
密封铅酸 (SLA) 电池	X	0	0	0	0	0
连接器/断路器触头	X	0	X	0	0	0

O: 表示此部件使用的所有同类材料中，此种有毒或有害物质的含量均低于 SJ/T11363-2006 规定的限制要求。

X: 表示此部件使用的至少一种同类材料中，此种有毒或有害物质的含量高于 SJ/T11363-2006 规定的限制要求。

此表中所有名称中含“X”的部件均符合欧盟 RoHS 立法——“2003 年 1 月 27 日欧洲议会和欧洲理事会关于在电气电子设备中限制使用某些有害物质的指令 2002/95/EC”。

只有使用荧光灯管的平板显示器中含有汞。只有不间断电源系统 (UPS) 和运行时间延长模块 (ERM) 中含有密封铅酸 (SLA) 电池。只有 UPS 和电源分配单元 (PDU) 中使用的连接器/断路器触头中可能含有镉。

注：所引用的环保使用期限标记根据产品的正常操作使用条件（如温度和湿度）确定。



除非另有标明，此电子信息产品污染控制标记适用于 HP 品牌的企业存储和服务器产品。

Index

Symbols

- \$SYSTEM, recovery operations for, 194
- \$YMIO.#CLCI, 201, 203
- \$YMIO.#CNSL, 201
- \$ZHOME, 202
- 4-Port ServerNet Extender (4PSE), 33, 35, 54, 108
- 4-Port ServerNet Extender (4PSE), servicing, 110

A

- Alternate system disk
 - creating, 184
- Asynchronous Terminal Process 6100 (ATP6100), 87
- ATP6100, 87, 215
- Automating
 - system shutdown, 201
 - system startup, 201

B

- BACKCOPY utility, 222
- BACKUP utility
 - backing up configuration and operations files, 138
 - description of, 222
- Batteries
 - maintaining, 173
- Battery ride-through, 196
- Binary number system, 232
- Binary to decimal conversion, 232
- BladeCluster Solution, 100

C

- Cartridge tape, handling and storing, 220
- CIIN file
 - contents, 203
 - establishing, 203
 - file name, 203
 - initial location of, 201
 - modifying, 204
 - ownership, 203
 - security, 203, 204
 - specifying, 203
 - system behavior when absent, 204
- Cleaning enclosures, 219
- CLIM
 - IB, 33, 35, 86
 - IP, 33, 35
 - powering off, 112
 - problems with, 114
 - states, 114
 - Storage, 33, 35
 - Telco, 33, 35, 86
- CLIMCMD mechanism
 - description of, 222
- climconfig command
 - description of, 222
- Cluster Essentials, 58

- see *also* NonStop Cluster Essentials
- Cluster Performance Essentials
 - for monitoring system performance, 72
 - See NonStop Cluster Performance Essentials, 58
- Collector (spooler), checking status of, 162
- Command files
 - examples, 202
- Communications line, 96
- Communications Process subsystem (CP6100), 87
- Compaq TSM see TSM
- Configuration
 - files, 203
 - INITIAL_COMMAND_FILE, 203
 - TCP/IP stacks, 210
- CONFLIST file, 138
- CONFTEXT, 203
 - see *also* CONFTEXT file
- CONFTEXT file, 138
 - INITIAL_COMINT_INFILE, 203
 - INITIAL_COMMAND_FILE, 203
- Converting numbers, 232
- CP6100, 87
- CPU n has been dumped to dumpfile message, 136
- Creating automated system monitoring procedures
 - using automated HP manageability tools, 64
 - using TACL macros, routines, command files, 64

D

- DCOM, 154, 222
- Decimal number system, 232
- Decimal to binary conversion, 235
- Decimal to hexadecimal conversion, 236
- Decimal to octal conversion, 236
- Dedicated Service LAN, 86
- Direct-connect line
 - startup file, 214
- Disk Compression Program (DCOM), 154, 222
- Disk drives
 - common problems, 150, 151
 - description of, 140
 - LEDs, 68
 - monitoring, 143
 - recovery operations for, 151, 152
- Disk Space Analysis Program (DSAP), 222
- Distributed Systems Management/Tape Catalog (DSM/TC), 223
- DSAP, 222
- DSM/TC, 223
- Dumping a processor to disk, 134
- Dumps
 - completed message, 136
 - dump file
 - checking with FUP, 136
 - submitting to service provider, 139

E

- EMS Analyzer (EMSA), [222](#)
- EMS event messages
 - monitored by HP SIM, [77](#)
 - monitored by OSM Event Viewer, [77](#)
 - monitored by Web ViewPoint, [78](#)
 - tools for monitoring, [77](#)
- EMS event messages, monitoring , [77](#)
- EMSA, [222](#)
- EMSDIST
 - description of, [222](#)
 - using to monitor EMS event messages, [78](#)
- EMSLOG file, [138](#)
- Enclosures
 - cleaning, [219](#)
- Encryption, volume level, [33](#), [37](#)
- Enterprise Storage System, [107](#)
 - see also ESS
- ESS, [107](#)
- Event Management Service (EMS), [77](#)
- Examples
 - checking file size, [149](#)
 - checking status of PATHMON process, [169](#)
 - checking status of TMF, [167](#)
 - MEDIACOM STATUS TAPE DRIVE command, [158](#)
 - SCF STATUS ADAPTER command, [89](#)
 - SCF STATUS DISK command, [146](#)
 - SCF STATUS LIF command, [90](#)
 - SCF STATUS LINE command, [97](#)
 - SCF STATUS PIF command, [90](#)
 - SCF STATUS TAPE command, [158](#)
- Expand-over-IP
 - startup file, [213](#)

F

- FCDM, [33](#), [35](#)
- FCSA, [86](#)
 - overview, [107](#)
 - problems with, [108](#)
 - states, [108](#)
- Fibre Channel disk module (FCDM), [33](#), [35](#)
- Fibre Channel ServerNet adapter, [106](#)
 - see also FCSA
- Fibre Channel ServerNet adapter (FCSA), [86](#)
- File Utility Program (FUP)
 - description of, [223](#)
 - INFO command, [136](#), [149](#)
- Freeze
 - enabling or disabling on a processor, [134](#)
 - freeze code=%nn message, [129](#)
 - hardware error, [128](#)
- FUP, [223](#)
 - see also File Utility Program (FUP)

G

- G4SA, [86](#)
 - monitoring, [109](#)
 - overview, [107](#)
 - states, [109](#)

- Gigabit Ethernet 4-port adapter (G4SA), [86](#)
- Guided procedures, OSM, [31](#)

H

- H-series, [13](#)
- Halting processors, [128](#), [129](#), [130](#)
- Hang
 - of processor, [128](#)
 - of system, recovery operations for , [130](#)
- Hexadecimal number system, [232](#)
- Hexadecimal to decimal conversion, [234](#)
- Home terminal, using \$ZHOME, [202](#)
- HP BladeSystem Integrated Manager, [56](#)
- HP BladeSystem Onboard Administrator (OA), used for system monitoring, [63](#)
- HP Insight Remote Support Advanced, [59](#), [223](#)
 - problem incident reports, [59](#)
- HP manageability tools
 - advantages of using them, [52](#)
 - for monitoring NonStop subsystem performance, [71](#)
 - for monitoring system resources, [52](#)
- HP NonStop Open System Management (OSM) see OSM
- HP NonStop Transaction Management Facility (TMF), [165](#)
 - see also TMF
- HP Operations Agent for NonStop (OVNM) see OVNM
- HP Performance Agent for NonStop (OVNPM) see OVNPM
- HP SIM
 - description of, [52](#), [55](#)
 - overview, [223](#)
 - plug-ins, [57](#)
- HP Systems Insight Manager see HP SIM
- HP Tandem Advanced Command Language (TACL), [139](#), [203](#)
 - see also TACL

I

- I/O adapter module (IOAM), [107](#)
- I/O Essentials, [57](#)
 - see also NonStop I/O Essentials
- InfiniBand CLuster I/O Module (IB CLIM), [33](#)
- INFO command (FUP), [136](#), [149](#)
- INITIAL_COMINT_INFILE, [203](#)
- INITIAL_COMMAND_FILE, [203](#)
- Insight Control Power Management
 - as HP SIM plug-in, [58](#)
 - for BladeSystems NB54000c, and NB54000c-cg, [58](#)
 - for BladeSystems NB56000c, and NB56000c-cg, [58](#)
- Insight Remote Support Advanced
 - See HP Insight Remote Support Advanced, [59](#)
- Integrated Lights-Out (iLO)
 - power cycling processors, [181](#)
- IOAM enclosure, [107](#)

J

- J-series, [13](#)

K

- Kernel-Managed Swap Facility (KMSF), [224](#)
- KMSF, [224](#)

L

LEDs

status, [68](#), [176](#)

LEDs, status, [68](#), [176](#)

LIFs, [86](#)

Logical interfaces (LIFs), [86](#)

M

Maintenance LAN, [86](#)

Measure program, [223](#)

MEDIACOM

description of, [223](#)

interface, [161](#), [223](#)

STATUS TAPEDRIVE command, [158](#)

Monitoring

communications subsystems, [98](#)

disk drives, [143](#)

EMS event messages, [77](#)

G4SA, [109](#)

NonStop system resources, [52](#)

overview, [51](#)

performance of NonStop subsystems, [71](#)

printers, [162](#)

problem incident reports, [59](#)

processes, [79](#)

processors, [124](#)

ServerNet fabrics, [102](#)

tape drives, [155](#)

N

NonStop Availability Statistics and Performance (ASAP)

for monitoring system performance, [72](#)

NonStop BladeSystem, [13](#)

NonStop BladeSystem NB50000c, [32](#), [33](#), [34](#)

NonStop BladeSystem NB50000c-cg, [32](#), [33](#), [34](#)

NonStop BladeSystem NB54000c, [32](#), [33](#), [34](#)

NonStop BladeSystem NB54000c-cg, [32](#), [33](#), [34](#)

NonStop BladeSystem NB56000c, [32](#), [33](#), [34](#)

NonStop BladeSystem NB56000c-cg, [32](#), [33](#), [34](#)

NonStop Cluster Essentials, [58](#), [223](#)

NonStop Cluster Performance Essentials, [58](#), [223](#)

for monitoring system performance, [72](#)

NonStop I/O Essentials, [57](#), [223](#)

NonStop Maintenance LAN DHCP DNS Configuration

Wizard, [224](#)

NonStop NET/MASTER application, [224](#)

NonStop NS-series, [13](#)

NonStop NS1000 system, [38](#)

NonStop NS1200 system, [38](#)

NonStop NS14000 series system, with VIO, [37](#)

NonStop NS14000 system, with IOAM, [37](#)

NonStop NS16000 series system, [37](#)

NonStop NS2000 series system, [36](#)

NonStop NS21000 system, [36](#)

NonStop NS2200 series system, [36](#)

NonStop NS23000 system, [36](#)

NonStop NS24000 series system, [36](#)

NonStop Software Essentials, [57](#), [224](#)

NonStop TCP/IP, [86](#)

NonStop TCP/IPv6, [86](#)

NonStop Virtual Hometerm Subsystem (VHS), [202](#)

NonStop Web ViewPoint

used to monitor NonStop subsystem performance, [73](#)

NonStop Web ViewPoint Storage Analyzer plug-in

for monitoring system performance, [74](#)

NSKCOM, [224](#)

Number conversion

binary to decimal, [232](#)

decimal to binary, [235](#)

decimal to hexadecimal, [236](#)

decimal to octal, [236](#)

hexadecimal to decimal, [234](#)

octal to decimal, [233](#)

overview, [232](#)

O

OA (HP BladeSystem Onboard Administrator), [224](#)

Octal number system, [232](#)

Octal to decimal conversion, [233](#)

Onboard Administrator (OA)

power cycling processors, [179](#)

used for system monitoring, [63](#)

OSM

CIIN file, [204](#)

description of, [224](#)

documentation, [228](#)

guided procedures, [31](#)

launching, [30](#)

security, [204](#)

OSM Event Viewer, [77](#)

OSM Service Connection

generating diagnostic data, [60](#)

generating diagnostic data for CLIMs, [88](#)

monitoring problem incident reports, [59](#)

recovery operations for problems detected by, [59](#)

used for system monitoring and serviceability, [59](#)

Outages, planned, [189](#)

OVNM

for monitoring system performance, [75](#)

for monitoring system resources, [64](#)

OVNPM

for monitoring system performance, [75](#)

for monitoring system resources, [64](#)

P

PAM, [86](#)

Parallel Library TCP/IP, [44](#)

PATHCOM, [167](#)

PATHMON processes, [167](#)

Pathway

commands, [167](#)

processes, [167](#)

transaction processing applications, [165](#), [167](#)

PEEK program, [224](#)

Physical interfaces (PIFs), [86](#)

PIFs, [86](#)

Planned outages, [189](#)

Port Access Method (PAM), [86](#)

POSTs, 176
 see also Power-on self-tests (POSTs)
Power failure
 how external devices respond to , 171
 preparing for
 maintaining batteries , 173
 monitor batteries , 173
 monitor power supplies, 173
 ride-through time , 173
 recovery operations, 174
 response
 ESS cabinets, 172
 external devices, 171
 IB CLIMs, 172
 IP CLIMs, 172
 NonStop cabinets, 171
 NonStop S-series enclosures, 171
 SAS disk drive enclosures, 172
 Storage CLIMs, 172
 systems , 171
 Telco CLIMs, 172

Powering off the system , 192

Powering on
 external system devices, 177
 power cycling processors, 179, 181
 system from low-power state, 176
 system from no power state, 177

Power-on self-tests (POSTs)
 system power-on, 176

Printers
 monitoring, 162
 recovery operations for, 162

Problems, common
 disk drive, 150, 151
 tape drive, 159

Processes
 generic, 80
 I/O, 79
 monitoring , 80
 recovery operations for, 84
 system, 79

Processor halts, 129
 halt code = %nn message, 129
 recovery operations for, 130

Processors
 dumping to disk, 134
 dumps, 135
 see also Dumps
 freeze, 129
 halt, 129
 halting processors, 130
 hang, 128
 monitoring, 59
 recovery operations for , 129

R

RCVDUMP utility, 136
Recovery operations
 for \$SYSTEM, 194

 for disk drives, 151, 152
 for printers, 162
 for processor halt, 130
 for processors , 128
 for ServerNet fabrics, 105
 for system console, 28
 for tape drives, 160
RESTORE utility, 225

S

SACs, 86
SCF, 225
 commands
 LISTDEV, 41
 STATUS ADAPTER, 89
 STATUS DISK, 145
 STATUS LIF, 90
 STATUS LINE, 97
 STATUS PIF, 90
 STATUS SAC, 89
 STATUS TAPE, 157
 STATUS, examples of, 61
 determining device states, 60
 managing disks, 161
 object states, 62
 powering off the system, 193
 storage device recovery, 151, 160
 used for system monitoring, 60
SCSI (SAS) disk drive enclosure, 33
ServerNet addressable controllers (SACs), 86
ServerNet clusters, 100
ServerNet connectivity, 100
ServerNet fabrics
 monitoring , 102
 recovery operations for, 105
ServerNet switch board, 107
Setting system time, 174
Shutdown files
 about, 214
 ATP6100 lines, 215
 automating, 201
 CP6100 lines, 215
 Expand-over-IP lines, 216
 security, 214
 sequence, 201, 214
 spooler, 217
 system shutdown file, 214
 TMF, 217
 X.25 lines, 216
SNAX/APN, 87
Software Essentials, 57, 224
 see also NonStop Software Essentials
Solid state disks (SSDs), 142, 172
SPOOLCOM, 225
Spooler, 209, 217
Startup files
 about, 203
 ATP6100 lines, 212
 automating, 201

- CIIN, 201
- configuration database, 208
- CP6100, 212
- direct-connect, 214
- Expand-over-IP, 213
- invoking, 201
- security, 207
- sequence, 208
- spooler warm start, 209
- system startup file, 208
- TCP/IP stacks, 208
- TMF warm start, 210
- X.25 lines, 213
- States
 - CLIM, 114
 - FCSA, 108
 - G4SA, 109
- Status LEDs, used to monitor system, 68
- Stopping the system, 191, 192
- Storing cartridge tapes, 220
- Subsystems
 - CIP, 43
 - displaying configuration of, 43
 - Kernel, 44
 - SLSA, 46, 85
 - storage, 44
 - TCP/IP, 44
 - WAN, 46, 85
- SWAN concentrator, 210
- System
 - performance, 190
 - powering off, 192
 - recording configuration of, 39
 - stopping, 191, 192
- System console
 - TACL window, 203
- System console, recovery operations for, 28
- System disk
 - creating alternate, 184
- System time, setting, 174

T

- TACL, 139, 203, 225
- Tape drives
 - common problems, 159
 - monitoring, 155
 - recovery operations for, 160
- Tapes, handling and storing, 220
- TCP/IP
 - configuration file, 210
 - startup file, 210
- TCP/IPv6, 86
- Terminals: monitoring; Monitoring: terminals; Terminals:
 - recovery operations for, 162
- TFDS, 134, 225
- TFDSCOM, 134, 225
- Time, system, setting, 174
- TMF, 210, 217
 - states of, 167

- STATUS command, 165, 166
- TMFCOM command, 165
- TMFCOM, 226
- TSM
 - CIIN file, 204
 - launching, 30
 - security, 204
- TSM workstation, 28
 - see also System console

V

- ViewSys utility, 128, 226
- VIO enclosure, description, 37
- VIO enclosure, powering on, 177
- Virtual Hometerm Subsystem (VHS);Hometerm, 202
 - see also VHS
- Volume level encryption, 33, 37

W

- Web ViewPoint
 - description of, 226
 - used to monitor EMS event messages, 78
 - used to monitor NonStop subsystem performance, 73

X

- X.25 lines, 216