IBM TotalStorage SAN File System
(based on IBM Storage Tank™ technology)

Maintenance and Problem Determination Guide

Version 1 Release 1

GA27-4318-00

IBM TotalStorage SAN File System
(based on IBM Storage Tank™ technology)

Maintenance and Problem Determination Guide

Version 1 Release 1

GA27-4318-00

NOTE
Before using this information and the product it supports, read the general information in Appendix H, "Notices", on page 413.

First Edition (November 2003)

This edition applies to the IBM TotalStorage SAN File System and to all subsequent releases and modifications until otherwise indicated in new editions.

Order publications through your IBM representative or the IBM branch office servicing your locality. Publications are not stocked at the address below.

IBM welcomes your comments. A form for reader's comments is provided at the back of this publication. If the form has been removed, you may address your comments to:

International Business Machines Corporation
Design & Information Development
Department CGFA
PO Box 12195
Research Triangle Park, NC 27709–9990
U.S.A.

You can also submit comments by selecting **Feedback** at www.ibm.com/storage/support.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Contents

About this guide vii
Who should use this guide viii
Notices in this book viii
Publications viii
Web sites x

Chapter 1. Getting started 1

Chapter 2. Introduction to SAN File System 3
Components 3
Engines 5
Administrative server 8
Metadata server 8
Clients 10
Master console 16
Service Alert 16
Remote access 17
SNMP 18

Chapter 3. Accessing SAN File System components 19
Remotely accessing the master console 20
Accessing the Administrative server through a browser 22
Accessing an engine through SSH 23
Accessing the RSA II adapter 24
Accessing a client through SSH. 25
Accessing a client through telnet 26
Accessing a client through a remote console utility 27

Chapter 4. Diagnostic tools 29
Engine diagnostic tools 29
POST 30
Light path diagnostics 30
Diagnostic tests 31
Engine error logs 32
OS Present diagnostics. 33
RSA II Web interface 35
Server diagnostic tools. 36
Logs. 36
Metadata server tracing 40
SAN File System dump capability 41
Linux dump capability 41
Client diagnostic tools 42
AIX client logging and tracing 42
AIX client dump capability 43
Windows 2000 client logging and tracing . . . 44
Windows 2000 client dump capability 45
One-button data collection utility 45
Hardware Vital Product Data 46
Software Vital Product Data 47

Chapter 5. Isolating problems with the SAN File System 51

Chapter 6. Troubleshooting the cluster 55
Troubleshooting a Metadata server 56
Troubleshooting the local network 59
Resolution Procedures 62
Shutting down an engine from the RSA Web interface 62
Taking a Metadata server offline 63
Defining a new master Metadata server 63
Reassigning filesets to Metadata servers 64
Bringing a Metadata server online 64
Repairing metadata. 64
Backing up system metadata 65

Chapter 7. Troubleshooting an Administrative server 67
Troubleshooting user access to the console 68
Troubleshooting user access to the Administrative command-line interface 70
Troubleshooting user access to a specific task or command 71
Resolution procedures 72
Verifying that the Administrative agent is running. 72
Verifying that the console is running 73
Replacing expired LDAP and CIMOM certificates 73
Configuring LDAP for SAN File System 74

Chapter 8. Troubleshooting a SAN File System client 77
Troubleshooting client access to data 77
Troubleshooting client performance problems . . . 80

Chapter 9. Troubleshooting the master console 83
Resolution procedures 83
Performing a total software recovery 83
Recovering a hard disk drive 84
Replacing fiber-channel cable and GBICs . . . 85

Chapter 10. Troubleshooting an engine 87
Performing general engine checkout 87
Resolving hardware component problems 88
Troubleshooting CD-ROM drive problems . . . 88
Troubleshooting diskette drive problems. . . . 89
Troubleshooting the Ethernet controller 90
Troubleshooting hard disk drive problems . . . 92
Troubleshooting keyboard, mouse, or pointing device problems 93
Troubleshooting memory problems 94
Troubleshooting microprocessor problems . . . 94

Troubleshooting monitor problems 95
Troubleshooting power problems 97
Troubleshooting USB port problems 99
Troubleshooting power-on problems. 100
Isolating undetermined problems. 101
Recovering BIOS code 101
Running engine diagnostics 104
Running OS Present diagnostics 105
Logging on to the RSA II 106
Starting the Configuration/Setup Utility 106

Chapter 11. Adding and replacing engine components 109
Removing the cover and bezel. 110
Adding or replacing an adapter 111
Replacing the battery 117
Replacing a CD-ROM drive. 119
Replacing the hard disk backplane 120
Replacing a diskette drive 121
Replacing a hot-swap hard disk drive 122
Replacing a hot-swap fan 123
Replacing a hot-swap power supply. 125
Replacing a memory module 127
Replacing a microprocessor. 131
Replacing the diagnostics/operator panel card . . 136
Replacing the power-supply cage. 137
Replacing an RSA II adapter 139
Replacing the system planar and shuttle 142
Replacing the cover and bezel 144

Chapter 12. Disaster recovery 145
Restoring the master console 145
Restoring the engine hardware and operating system 146
Restoring SAN connectivity 146
Restoring SAN File System software. 147
Restoring SAN File System cluster configuration 147
Restoring SAN File System metadata 149
Restoring SAN File System clients 150
Restoring SAN File System user data 150

Chapter 13. Getting help, service, and information 151
Getting help online 152
Before you call for service 152
Getting help by telephone 152

Appendix A. Accessibility 153

Appendix B. Cabling diagram 155

Appendix C. Commands. 157
Administrative commands 158
activatevol 163
addsnmpmgr 164
attachcontainer 166
builddrscript 168
catlog 169
catpolicy 171
chclusterconfig 172
chcontainer 174
chpool. 176
chvol 178
clearlog 180
collectdiag 181
detachcontainer 183
exit. 185
help 186
lsadmuser 187
lsautorestart 190
lsclient 194
lscontainer 198
lsdrfile 202
lsengine 204
lsimage 207
lslun 210
lspolicy 213
lspool 216
lsproc 219
lsserver 221
lssnmpmgr 225
lstrapsetting 226
lsvol 227
mkcontainer 230
mkdrfile 233
mkimage 234
mkpolicy 236
mkpool 238
mkvol 240
quiescecluster 242
quit 244
reportvolfiles 245
resetadmuser 246
restartengine 247
resumecluster 248
reverttoimage 249
rmcontainer 251
rmdrfile 252
rmimage 253
rmpolicy 255
rmpool 256
rmsnmpmgr 257
rmvol 258
setcontainerserver 260
setdefaultpool 262
setmaster 263
setoutput 266
settrap. 268
startautorestart 270
startcluster 272
startengine 273
startmetadatacheck 274
startserver 276
statcluster 278
statcontainer. 281
statengine 282
statserver. 286
stopautorestart 289
stopcluster 290
stopengine 291

stopmetadatacheck 292
stopserver 293
suspendvol 294
upgradecluster 295
usepolicy 296
Client commands 297
AIX-client commands. 297
Windows-client command 315
Service commands and utilities 319
addserver 322
chadvclusterconfig. 323
disableConsoleTrace 325
dropserver 326
enableConsoleTrace 327
legacy 328
mktruststore. 335
pmf 337
startCimom 343
startConsole 344
stfsdebug. 345
stfsstat. 347
stopCimom 351
stopConsole 352
tank extractbootrecord 353
tank lscluster 355
tank lsversion 357
tank lsdisklabel. 358
tank resetcluster 360
tank resetversion 362

Appendix D. Engines 363
Major components of the Model 1RX 363
Controls and indicators 365
System board 367
Input and output ports 372

Appendix E. Hardware messages and codes. 377
Beep symptoms 377
Bus fault error messages. 380
Hard disk drive error messages 381
Diagnostic error codes 381
Fan error messages 386
Host built-in self test (BIST) error messages . . . 387
Light path diagnostic panel LEDs 387
POST error codes 389
Power error messages 397
Temperature error messages 399
Temperature-related system shutdown messages 400
Voltage-related system shutdown messages . . . 402

Appendix F. Parts listings 405
Model 1RX 405
Power cords. 407

Appendix G. Customer-replaceable components 411
Replacing storage engine components 411
Replacing master console components 411

Appendix H. Notices 413
Electronic emission notices 414
Federal Communications Commission (FCC) statement. 414
Industry Canada Class A emission compliance statement. 415
Australia and New Zealand Class A statement 415
United Kingdom telecommunications safety requirement 415
European Union EMC Directive conformance statement. 415
Germany Electromagnetic Compatibility Directive 416
Chinese Class A warning statement 417
Taiwan electrical emission statement. 417
Japanese Voluntary Control Council for Interference (VCCI) statement 417
End of life statement 417
Safety information. 418
Basic safety information (multilingual translations) 418
General safety 422
Grounding (earthing) requirements 423
Handling electrostatic discharge-sensitive devices 423
Handling static-sensitive devices 424
Safety inspection guide 424
System reliability considerations 425
Working inside the engine with the power on 426
Before you begin replacing components . . . 426
Trademarks 427

Glossary 429

Index 433

About this guide

This guide provides maintenance and problem determination information about the IBM TotalStorage® SAN File System hardware and software. The information is organized as follows:

- Chapter 1, “Getting started”, on page 1 explains how to begin problem determination for the SAN File System.
- Chapter 2, “Introduction to SAN File System”, on page 3 provides an overview of SAN File System and related concepts.
- Chapter 3, “Accessing SAN File System components”, on page 19 describes the various methods that are available for accessing SAN File System components, including hardware and software.
- Chapter 4, “Diagnostic tools”, on page 29 provides information about the tools that you can use to diagnose problems with SAN File System components.
- Chapter 5, “Isolating problems with the SAN File System”, on page 51 describes initial steps that you can take to begin isolating problems with the SAN File System, IP network, and SAN.
- Chapter 6, “Troubleshooting the cluster”, on page 55 explains how to diagnose and resolve problems related to the SAN File System cluster, including the Metadata servers and the IP network. In addition, it provides procedures that can assist you in resolving problems with the cluster.
- Chapter 7, “Troubleshooting an Administrative server”, on page 67 explains how to diagnose and resolve problems related to the Administrative server. It includes information about problems related to administrative access to the SAN File System console, and the Administrative command-line interface. In addition, it provides procedures that can assist you in resolving problems with the Administrative server.
- Chapter 8, “Troubleshooting a SAN File System client”, on page 77 explains how to diagnose and resolve problems related to client access to user data as well as client performance. It also provides procedures that can assist you in resolving problems with clients.
- Chapter 9, “Troubleshooting the master console”, on page 83 explains how to diagnose and resolve problems related to the master console.
- Chapter 10, “Troubleshooting an engine”, on page 87 explains how to diagnose and resolve problems related to the engine.
- Chapter 11, “Adding and replacing engine components”, on page 109 explains how to replace the hardware components of an engine.
- Chapter 12, “Disaster recovery”, on page 145 provides the procedures for recovering the SAN File System cluster.
- Chapter 13, “Getting help, service, and information”, on page 151 explains how to obtain help.
- The appendices provide the following additional information:
  - Accessibility features of the SAN File System console and help system
  - Cabling diagram for the engines in the cluster
  - Commands for the Administrative command-line interface, SAN File System clients, and service.
  - Parts diagrams for engine and its components.
  - Messages and codes related to hardware.

– A list of the parts for the SAN File System.
– A list of customer-replaceable components.
– Notices

Who should use this guide

This guide is intended primarily for hardware and software service personnel who are familiar with the SAN File System and the IBM TotalStorage 4146 Model 1RX.

- The hardware diagnostic procedures and most replacement procedures should only be used by qualified support personnel. The only exception to the replacement procedures are the customer replaceable components listed in the appendix of this document.
- The service commands and utilities listed in this document may cause a disruption in the SAN File System. Therefore, these commands should only be used by qualified software support personnel.

Notices in this book

The following notices are contained within this guide and convey these specific meanings:

Note: These notices provide important tips, guidance, or advice.

Attention: These notices indicate possible damage to programs, devices, or data. An attention notice appears before the instruction or situation in which damage could occur.

CAUTION:
These notices indicate situations that can be potentially hazardous to you. A caution notice appears before the description of a potentially hazardous procedure step or situation.

DANGER

These notices indicate situations that can be potentially lethal or extremely hazardous to you. A danger notice appears before a description of a potentially lethal or extremely hazardous procedure step or situation.

Publications

The following publications are available in the SAN File System library. They are provided in softcopy on the *IBM TotalStorage SAN File System Publications CD* that came with your storage engine and at www.ibm.com/storage/support. To use the CD, insert it in the CD-ROM drive. If the CD does not launch automatically, follow the instructions on the CD label.

Note: The softcopy version of these publications are accessibility-enabled for the IBM Home Page Reader.

- *IBM TotalStorage SAN File System Release Notes*

  This document provides any changes that were not available at the time the publications were produced. This document is available only from the technical support Web site: www.ibm.com/storage/support
- *IBM Safety Information — Read This First*, SD21-0030

This document provides translated versions of general safety notices and should be read before using this product. This document is provided only in hardcopy.

- *IBM Statement of Limited Warranty*

  This publication describes the IBM statement of limited warranty as it applies to the SAN File System Model 1RX storage engine.

- *IBM TotalStorage SAN File System Software License Information*

  This publication provides multilingual information regarding the software license for IBM TotalStorage SAN File System Software.

- *IBM TotalStorage SAN File System Administrator's Guide and Reference*, GA27-4317

  This publication introduces the concept of SAN File System, and provides instructions for configuring, managing, and monitoring the system using the SAN File System console and Administrative command-line interfaces. This book also contains a commands reference for tasks that can be performed at the administrative and client command-line interfaces.

- *IBM TotalStorage SAN File System Maintenance and Problem Determination Guide*, GA27-4318

  This publication provides instructions for adding and replacing hardware components, monitoring and troubleshooting the system, and resolving hardware and software problems.

  Note: This document is intended only for trained personnel.

- *IBM TotalStorage SAN File System Messages Reference*, GC30-4076

  This publication contains message description and resolution information for errors that can occur in SAN File System software.

- *IBM TotalStorage SAN File System Planning, Installation and Configuration Guide*, GA27-4316

  This publication provides detailed procedures to plan the installation and configuration of SAN File System, set up and cable the hardware, perform the minimum required configuration, migrate existing data, and upgrading software.

- *Rack Installation Instructions*

  This publication provides instructions for installing the Model 1RX in a rack.

- *IBM TotalStorage SAN File System System Management API Guide and Reference*, GA27-4315

  This publication contains guide and reference information for using the CIM Proxy API, including common and SAN File System-specific information.

  Note: This document contains information and procedures intended for only selected IBM Business Partners. Contact your IBM representative before using this publication.

- *Subsystem Device Driver User's Guide for the IBM TotalStorage Enterprise Storage Server and the IBM TotalStorage SAN Volume Controller*, SC26-7540

  The Subsystem Device Driver (SDD) provides the multipath configuration environment support for a host system that is attached to an IBM TotalStorage Enterprise Storage Server® (ESS), IBM TotalStorage SAN Volume Controller or IBM TotalStorage SAN File System. This book provides step-by-step procedures on how to install, configure, and use SDD for the host systems.

  Note: SAN File System supports the version of the Subsystem Device Driver that is shipped with the program product.

- *IBM TotalStorage Translated Safety Notices*, GA67-0043

This publication contains translated versions of hardware caution and danger statements that appear in the publications in this library. Each caution and danger statement has an assigned number that you can use to locate the corresponding statement in your native language.

Web sites

The following Web sites have additional and up-to-date information about SAN File System:

- www.ibm.com/storage/support
- www.ibm.com/storage/software/virtualization/sfs

Chapter 1. Getting started

This topic is the starting point for all SAN File System problem determination actions. Using this topic, service representatives can quickly determine the appropriate chapter in this guide for their particular maintenance action. It guides the service representative through the process of determining the location of a failure, of preparing the SAN for maintenance, and of performing the repair.

Service representatives should perform the following steps:

1. Attach a keyboard and display to the SAN File System engine.

   You can attach a keyboard and display to the specific engine or you can access the engine from the master console. See Chapter 3, "Accessing SAN File System components", on page 19 for the methods that you can use to access engine.

   Note: Depending on the proximity of the master console to the engines in the SAN File System cluster, you may need to locally attach a customer-supplied keyboard, monitor, and mouse to the engine before attempting to service the engine.

2. Determine whether the problem is within the SAN File System subsystem.

   To determine whether the problem is within SAN File System or elsewhere in the SAN, see Chapter 5, "Isolating problems with the SAN File System", on page 51 to make the determination.

   If you determine that the problem is not within the SAN File System, follow the instructions in Chapter 5, "Isolating problems with the SAN File System", on page 51 to resolve the problem.

3. Determine whether the problem is within a SAN File System Metadata server engine.

   To determine whether the problem is within a Metadata server engine, see Chapter 6, "Troubleshooting the cluster", on page 55.

   - If the problem is determined to be in the SAN File System master console, see Chapter 9, "Troubleshooting the master console", on page 83.
   - If the problem is determined to be neither the master console nor an engine, call your next level of support for assistance.

4. Determine which SAN File System engine is the cause of the problem.

   To determine which Metadata server engine is causing the problem, see Chapter 10, "Troubleshooting an engine", on page 87.

5. Request that the customer prepare the engine for maintenance.

   Before you begin working on an engine, it must be taken offline from the SAN File System cluster. Verify with the customer that the engine has been taken offline before attempting to troubleshoot or replace any hardware components.

6. Perform the maintenance action.

   See Chapter 11, "Adding and replacing engine components", on page 109 for information about replacing hardware components in the engine. After the engine has been repaired, request that the customer bring the engine back online within the SAN File System cluster.

Chapter 2. Introduction to SAN File System

IBM TotalStorage SAN File System is a multiplatform, scalable file system and storage management solution that works with a storage area network (SAN). It uses SAN technology, which allows an enterprise to connect large numbers of devices, such as client and server machines and mass storage subsystems, to a high-performance network.

On a SAN, storage is separated from the computers that use it. With SAN File System, heterogeneous clients can access shared data directly from large, high-performance, high-function storage systems, such as IBM TotalStorage Enterprise Storage Server®. The SAN File System is currently built on a fibre-channel network, and is designed to provide superior I/O performance for data sharing among heterogeneous computers. It also provides growth capability and simplified storage management.

SAN File System differs from conventional distributed file systems in that it uses a data-access model that requires *clients* to contact servers to obtain only the information they need to locate and access data on *storage devices*, not the data itself. SAN File System clients access data directly from storage devices using the high bandwidth provided by a fibre-channel network. Direct data access eliminates server bottlenecks and provides the performance necessary for data-intensive applications.

SAN File System presents a single, *global namespace* to clients where they can create and share data. Data consistency and integrity is maintained through SAN File System's management of distributed *locks* and the use of *leases*. SAN File System provides locks that enable file sharing among SAN File System clients, and when necessary, provides locks that allow clients to have exclusive access to files. A lease determines the maximum period of time that a server guarantees the locks it grants to clients. A client must contact the server before the lease period ends in order to retain its locks.

SAN File System also provides the benefits of automatic *file placement* through the use of *policies* and *rules*. Based on rules specified in a policy set, SAN File System automatically stores data on devices in *storage pools* that are specifically created to provide the capabilities and performance appropriate for how the data is accessed and used.

Components

A *component* is one of the elements that make up the entire SAN File System. The following figure illustrates the major SAN File System software components.

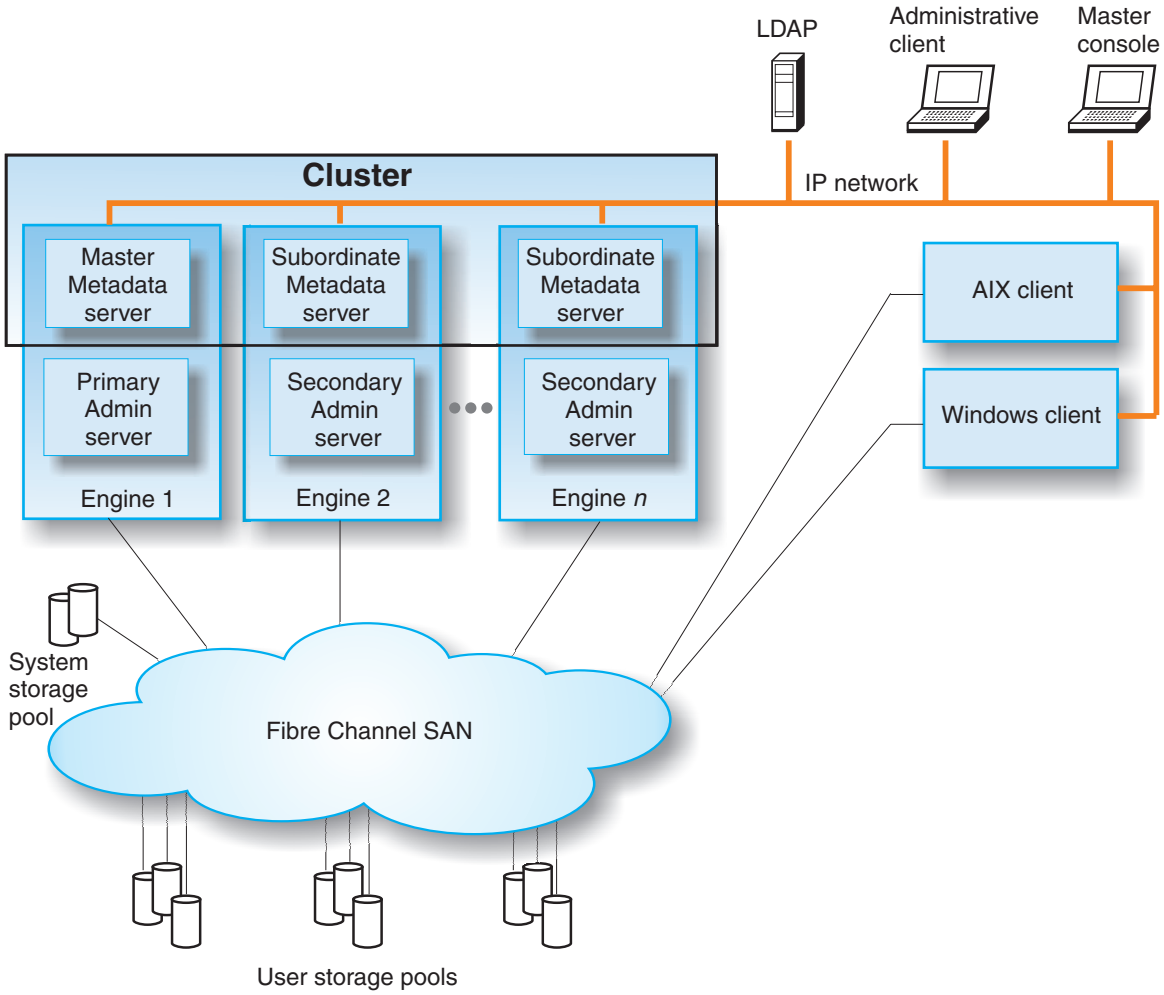

Figure 1. SAN File System components

The Metadata servers and clients communicate over an internal IP network and access data over the SAN.

Note: SAN File System relies on preexisting networking hardware, including an IP network, SAN, network switches, and routers.

A *Metadata server* is a server that runs on a SAN File System engine and performs metadata, administrative, and storage management services. Multiple Metadata servers form a server *cluster*. In a SAN File System server cluster, there is one master Metadata server and one or more subordinate Metadata servers, each running on a separate engine in the cluster. Together these Metadata servers provide clients with shared, coherent access to the SAN File System *global namespace*.

The *metadata* resides on private storage that is shared among all the Metadata servers in the cluster.

The *Administrative server* allows SAN File System to be remotely monitored and controlled through a Web-based user interface, called the *SAN File System console*, using either Netscape 6.2 and later or Internet Explorer 6.0 and later. The Administrative server also processes requests issued from the administrative command line interface, which is called *tanktool*. The Administrative server uses an

LDAP server, also connected to the internal IP network, to look up authentication and authorization information about the administrative users. The primary Administrative server runs on the same engine as the master Metadata server. It receives all requests issued by administrators and also communicates with the Administrative servers that run on each additional server in the cluster to perform routine requests.

The *clients* access the global namespace through a virtual or installable file system, which is installed on the client machine. These clients can act as servers to a broader clientele, providing Network File System (NFS) or Common Internet File System (CIFS) access to the global namespace or hosting applications (such as database servers or Web-hosting services that use multiple servers).

Related topics:

- "Administrative server" on page 8
- "Clients" on page 10
- "Engines"
- "Metadata server" on page 8

Engines

Within SAN File System, an *engine* is the hardware (which is based on the IBM xSeries® platform) on which a Metadata server and an Administrative server run. Each engine comes preloaded with SAN File System server software and must be configured by onsite IBM personnel. SAN File System supports from two to eight Model 1RX engines.

Note: Although you cannot purchase SAN File System with only one engine, you can run a single-engine system if all of the other engines fail (for example, if you have only two engines, and one of them fails), or if you want to bring down all of the engines except one before performing scheduled maintenance tasks.

The administrative infrastructure on each engine allows an administrator to monitor and control SAN File System from a standard Web browser or an administrative command line interface. The two major components of the infrastructure are an Administrative agent, which provides access to administrative operations, and a Web server that is bundled with the console services and servlets that render HTML for the administrative browsers. The infrastructure also includes a Service Location Protocol (SLP) daemon, which is used for administrative discovery of SAN File System resources by third-party Common Information Model (CIM) agents, and an IBM Director agent, which detects when the operating system "hangs" or stops unexpectedly and can reboot the operating system.

An administrator can use the SAN File System console, which is the Web-based user interface, or administrative commands to monitor and control an engine.

Related topics:

- "Hardware features" on page 7
- "Hardware specifications" on page 6
- "Metadata server" on page 8
- "Major components of the Model 1RX" on page 363

Hardware specifications

Dimensions:

Height	85.4 mm (3.36 in.)
Depth	698 mm (27.48 in.)
Width	443.6 mm (17.5 in.)
Weight	21.09 kg to 28.12 kg (46.5 lb to 62 lb)

Operating environment:

Air temperature (at maximum altitude of 2134 m [7000 ft])	10° to 35°C (50° to 95°F)
Humidity	8% to 80%

Nonoperating (power-off) environment:

Air temperature (at maximum altitude of 2134 m [7000 ft])	–40° to 60°C (–40°0 to 140°F)
Humidity	8% to 80%

Heat output:

Minimum configuration	341 Btu/hour (100 watt)
Maximum configuration	2200 Btu/hour (645.2 watt)

Power requirements:

Voltage low range	90 V ac to 137 V ac
Voltage high range	180 V ac to 265 V ac
Power consumption	0.1 kVA to 0.62 kVa
Sine-wave input	50 to 60 Hz required

Acoustical noise emission:

Declared sound power	6.7 bel
Bystander sound pressure	50 dBa

Related topics:

- "Hardware features" on page 7
- "Hardware requirements"

Hardware requirements

Customers are required to provide the following hardware:

- Internal IP network and Ethernet cables
- Fibre Channel SAN and fiber-channel cables
- Rack cabinet
- Ethernet hub or switch that will be used to connect the RSA II adapter on each engine with the master console. The Ethernet hub or switch must have sufficient ports to accommodate all SAN File System engines in the cluster and the master console.

- At least one Gigabit Ethernet cable (fiber) for each engine
- Two power outlets, one for each power supply, preferably on separate circuits.

Racks:

The Model 1RX is housed in a 19-inch, 2U tall, rack-mountable drawer and can be installed in a rack that meets the following requirements:

- Minimum depth of 70 mm (2.76 in.) between the front-mounting flange and the inside of the front door
- Minimum depth of 157 mm (6.18 in.) between the rear-mounting flange and the inside of the rear door
- Minimum depth of 718 mm (28.27. in.) and a maximum depth of 762 (30 in.) between the front- and rear-mounting flanges to support the use of the cable-management arm
- Cabinet must have perforated doors

Power outlets:

The power outlets must adhere to one of the following voltage ratings:

- 5 amps at 100 - 127 V ac at 50/60 Hz single phase
- 2.5 amps at 200 - 240 V ac at 50/60 Hz single phase

Related topics:

- "Hardware features"
- "Hardware specifications" on page 6
- "Components" on page 3

Hardware features

Standard features:

Hardware	Feature
Engine	• Two 2-way 3.067-GHz processor • 512-KB Level-2 cache • Two integrated 10/100/1000 Ethernet ports • IBM integrated system management processor (ISMP) • Five PCI expansion slots
Memory	• 4 GB • Type: 100-MHz, ECC, double-data rate (DDR), SDRAM, registered DIMMs
Drives	• 1.44-MB diskette • CD-ROM • Two hot-swap RAID drives (one drive is a mirror of the other)
Power supplies	Two 514-watt (115–230 V ac) hot-swap power supplies
Cooling	Eight redundant fans
Adapters	• One 2-port Gigabit (1 and 2-Gb/s) fibre-channel adapter • IBM Remote Supervisor Adapter II (RSA II)

Optional features:

Hardware	Feature
Adapters	Two 1-port Gigabit fiber-channel Ethernet adapters

Related topics:

- "Hardware requirements" on page 6
- "Hardware specifications" on page 6

Administrative server

The SAN File System *Administrative server*, which is based on a Web server software platform, is made up of the following parts:

- The GUI Web server, which renders the Web pages that make up the SAN File System console. The console is a Web-based user interface. It can be accessed using any standard Web browser, such as Mozilla or Microsoft® Internet Explorer, that has network access to the engines that host the Metadata servers in a server cluster.
- The Administrative agent, which implements all of the management logic for the GUI, CLI, and CIM interfaces, as well as performing administrative authorization and authentication against LDAP. The Administrative agent processes all management requests initiated by an administrator from the SAN File System console, as well as requests initiated from the administrative command line interface, which is called *tanktool*. The Agent communicates with the SAN File System Metadata servers, the OS, the RSA II card, the LDAP, and Administrative agents on other nodes in the cluster when processing requests.

An Administrative server interacts with a Metadata server through the Administrative agent. When you issue a request, the Administrative agent checks with an LDAP server, which must be installed in your environment, to authenticate the user ID and password and to verify whether the user has the authority (is assigned the appropriate role) to issue a particular request. After authenticating a user, the Administrative agent interacts with the Metadata server on behalf of that user to process the request. This same system of authentication and interaction is also available to third-party CIM clients to manage SAN File System.

To ensure high availability, an Administrative server resides on each engine of a cluster. All requests are received by the Administrative server that runs on the same engine as the master Metadata server. This is the primary Administrative server. However, requests can also be processed by Administrative servers running on other engines, and all requests are redirected to an Administrative server running on another engine if the primary Administrative server is not available.

Related topics:

- "Engines" on page 5
- "Metadata server"

Metadata server

A *Metadata server* is a server that runs on a SAN File System engine and performs metadata, administrative, and storage management services. In a SAN File System server cluster, there is one master Metadata server and one or more subordinate

Metadata servers, each running on a separate engine in the cluster. Together these Metadata servers provide clients with shared, coherent access to the SAN File System global namespace.

All of the servers, including the master Metadata server, share the workload of the SAN File System global namespace. Each is responsible for providing metadata and locks to clients for specific filesets assigned to them by an administrator. They know which filesets belong to which server, and when contacted by a client can direct the client to the appropriate server. They manage distributed locks to ensure the integrity of all of the data within the global namespace.

In addition to providing metadata to clients and managing locks, Metadata servers perform a wide variety of other tasks. They process requests issued by administrators to create and manage filesets, storage pools, volumes, and policy sets, and they enforce the policies defined by administrators to place files in appropriate storage pools and ensure that capacity quotas established for filesets and storage pools are not exceeded.

Performing Metadata service:

There are two types of metadata:

- *File metadata*, which is information that clients need to access files directly from storage devices on your storage area network. File metadata includes permissions, owner and group, access time, creation time, and other file characteristics.
- *System metadata*, which is metadata used by the system itself. System metadata includes information about filesets, storage pools, volumes, and policies. It is stored and managed in a separate system storage pool that is only accessible by the Metadata servers in your server cluster. The Metadata servers perform the reads and writes required to create, distribute, and manage this information.

Distributing locks to clients involves the following:

- Issuing leases that determine the length of time that a server guarantees the locks that it grants to clients.
- Granting locks to clients that allow them shared or exclusive access to files or parts of files. These locks are semi-preemptible, which means that if a client does not contact the server within the lease period, the server can “steal” the client’s locks and grant them to other clients if requested; otherwise, the client can reassert its locks (get its locks back) when it can contact the server again.
- Providing a grace period during which a client can reassert its locks before other clients can obtain new locks if the server itself goes down and then comes back online.

Performing administrative services:

A Metadata server processes requests from administrators (issued from the SAN File System console or by using administrative commands) to perform the following types of tasks:

- Create and manage filesets, which are subsets of the entire global namespace and serve as the units of workload assigned to specific Metadata servers.
- Create and manage volumes, which are LUNs labeled for SAN File System’s use in storage pools.
- Create and maintain storage pools (for example, an administrator can create a storage pool that consists of RAID or striped storage devices to meet reliability

requirements, and can create a storage pool that consists of random-access or low-latency storage devices to meet high performance requirements).
- Create FlashCopy images of filesets in the global namespace that can be used to make file-based backups easier to perform.
- Define policy sets that contain rules that determine in which storage pools specific files are stored.

Performing storage management services:

A Metadata server performs these storage management services:
- Manages allocation of blocks of space for files on LUNs
- Maintains pointers to the bits of a file that constitute the file
- Evaluates the rules in the active policy set and manages the placement of files in specific storage pools based on those rules
- Issues alerts when filesets and storage pools reach or exceed their administrator-specified thresholds, or returns out-of-space messages if they run out of space

Using configuration parameters:

Each Metadata server has a configuration file that is stored on local storage for each engine in a Metadata server cluster. This configuration file contains settings, which are specified by an administrator, for the following:

Space reclamation interval
Controls how often the master Metadata server reclaims free storage pool partitions allocated by filesets.

Lease period
Specifies the maximum length of time that a server guarantees the validity of the locks it has granted to a client. The client must contact the server before the lease period ends to renew its lease and retain its locks.

List of privileged clients
Lists the clients on which users specified as root or administrator users can have those same privileges on the SAN File System global namespace.

List of SNMP managers
Specifies the IP address, port, version (SNMPv1 or SNMPv2c) and community strings for SNMP trap recipients.

SNMP trap setting
Specifies the filter used on the cluster message log to determine which messages (informational, warning, error, and severe) are also issued as SNMP trap messages.

Related topics:
- "Clients"
- "Engines" on page 5
- "SNMP" on page 18
- "Components" on page 3

Clients

SAN File System is based on a client-server design. A SAN File System *client* is a computer that accesses and creates data that is stored in the SAN File System

global namespace. SAN File System uses a cluster of Metadata servers that cooperate with each other to provide the global namespace.

The *SAN File System protocol specification* includes a description of the protocols used between a metadata server and clients running on application servers. It is available at www.ibm.com/storage/software/virtualization/sfs.

Clients access only metadata (such as a file's location on a storage device) from a Metadata server, and then access data directly from storage devices attached to your high-speed fibre-channel network. This method of data access eliminates server bottlenecks and provides read and write performance that is comparable to that of file systems built on bus-attached, high-performance storage.

SAN File System supports clients that run these operating systems:

- AIX® 5.1 (32–bit only)
- Windows® 2000 Server or Windows 2000 Advanced Server

Client software must be installed on each AIX or Windows client. On an AIX client, the software is a virtual file system (VFS), and on a Windows client, it is an installable file system (IFS). The VFS and IFS provide clients with local access to the global namespace on your SAN. Note that clients can also act as servers to a broader clientele. They can provide NFS or CIFS access to the global namespace and can host applications such as database servers.

A VFS is a subsystem of an AIX client's virtual file system layer, and an IFS is a subsystem of a Windows client's file system. A VFS or IFS directs all metadata operations to a Metadata server and all data operations to storage devices attached to your SAN. A VFS or IFS makes the metadata that is visible to a client's operating system, as well as any applications that the client runs, look identical to metadata read from a native, locally attached file system.

When the global namespace is mounted on an AIX client, it looks like a local file system. When the global namespace is mounted on a Windows client, it appears as another drive letter and looks very much like an NTFS file system. Basically, the SAN File System global namespace looks and acts like any other file system on the client's system.

Caching metadata, locks, and data:

Caching allows a client to achieve low-latency access to both metadata and data. A client can cache metadata to perform multiple metadata reads locally. The metadata includes the mapping of logical file system data to physical addresses on storage devices attached to your SAN.

A client can also cache locks to allow the client to grant multiple opens to a file locally without having to contact a Metadata server for each operation that requires a lock.

In addition, a client can cache data for small files to eliminate I/O operations to storage devices attached to your SAN. A client performs all data caching in memory. Note that if there is not enough space in the client's cache for all of the data in a file, the client simply reads the data from the shared storage device on which the file is stored. Data access is still fast because the client has direct access to all storage devices attached to your SAN.

Using the direct I/O mode:

Some applications, such as database management systems, use their own sophisticated cache management systems. For such applications, SAN File System provides a direct I/O mode. In this mode, SAN File System performs direct writes to disk, does not cache data, and allows distributed applications on different computers to write data to the same file at the same time. Using the direct I/O mode makes files behave more like raw devices. This gives database systems direct control over their I/O operations, while still providing the advantages of SAN File System, such as the FlashCopy Image feature and file-level backup and restore.

Sharing files:

In a homogenous environment, which means that the clients in your environment are either all AIX or all Windows clients, SAN File System provides access and semantics that are customized for the operating system running on the clients. When files are created and accessed from only Windows clients, all the security features of Windows are available and enforced. When files are created and accessed from only AIX clients, all the security features of AIX are available and enforced.

In a heterogeneous environment, which means that you have both AIX and Windows clients, there is a restricted form of access. When files created on an AIX client are accessed by a Windows client, access is controlled using only the semantics and permissions of the "Other" permission bits in AIX. Similarly, when files created on a Windows client are accessed on an AIX client, access is controlled using only the semantics and permissions of the "Everyone" user group in Windows.

Antivirus software:

If more than one SAN File System client for Windows is running antivirus software that scans directories and files, there is no reason to run this from more than one SAN File System client on the same files. If scans are run from more than one client, they should be scheduled to run at different times to allow better performance of each scan.

Related topics:

- "Components" on page 3
- "UNIX-based clients"
- "Windows-based clients" on page 13

UNIX-based clients

A *UNIX®-based client* is a SAN File System client that runs a UNIX operating system and has the SAN File System client code installed. In this release, SAN File System supports clients running AIX 5.1 (32–bit only).

UNIX-based clients use standard UNIX permission semantics (such as read, write, and execute bits, and owner and group IDs) that make the SAN File System global namespace appear as if it is a local UNIX file system.

The SAN File System client code installed on a UNIX client is called a Virtual File System (VFS). The VFS is a subsystem of an AIX client's virtual file system layer. It directs all metadata operations to a SAN File System Metadata server and all data operations to storage devices attached to your SAN. A VFS makes the metadata

that is visible to a client's operating system, as well as any applications that run on the client, look identical to metadata read from a native, locally attached file system.

Privileged clients:

A privileged client is a client machine on which root users are given the same privileges on the global namespace as they have on other file systems available on their system. A root user logged in to a privileged AIX client is granted full control over directories, files, and other file system objects created by AIX clients.

If a root user logs in to a client that is not a privileged client, that user's privileges for the global namespace are reduced to those of "Other."

Commands:

There is a set of commands (issued from the operating system command line) that is used to set up, start and stop the client and to migrate data. These commands are accessible from each client machine. The client commands are separate from the administrative command line interface. The client commands allow users to perform operations on the client systems, and the administrative command line interface allows administrators to administer all aspects of the SAN File System server.

Constraints and restrictions:

The following are restrictions for AIX client features:

- AIX clients cannot use user IDs or group IDs 999999 and 1000000 for real users or groups; these are reserved IDs used internally by SAN File System.
- Multibyte enablement for items referenced in a file placement policy (storage pools, filesets, parts of file names) is not supported.
- mkfs and mkfifo (named pipes) in AIX are not supported.

Related topics:

- "Clients" on page 10
- "Components" on page 3
- "Windows-based clients"

Windows-based clients

A *Windows-based client* is a client that runs a Windows® operating system and has the SAN File System client code installed. In this release, SAN File System supports clients running on Windows 2000 Advanced Server.

The SAN File System client code installed on a Windows client is called an Installable File System (IFS). The IFS is a subsystem of a Windows client's file system. It directs all metadata operations to a SAN File System Metadata server and all data operations to storage devices attached to your storage area network (SAN). An IFS makes the metadata that is visible to a client's operating system, as well as any applications that run on the client, look identical to metadata read from a native, locally attached file system.

Windows-based clients use a subset of the Windows semantics. The supported semantics are described to Windows as volume properties, which are visible, for example, as properties of the drive within Windows Explorer. The following volume properties are supported by SAN File System:

- NTFS-like access control lists (requires all Windows-based clients to share a common Active Directory domain for users and groups)
- Long names and short names (eight-character names with three-character extensions)
- UNICODE-based file names
- Case-sensitive file names

Case sensitivity:

While the global namespace is case-sensitive, Windows applications can choose to use case sensitive or case insensitive names. This means that case-sensitive applications, such as those making use of Windows support for POSIX interfaces, behave as expected. Native Win32 clients (such as Windows Explorer) get only case-aware semantics. The case specified at time of file creation is preserved, but in general, file names are case-insensitive. For example, Windows Explorer allows the user to create a file named "Hello.c," but an attempt to create "hello.c" in the same folder will fail because the file already exists. If a Windows-based client accesses a folder that contains two files that are created on a UNIX-based client with names that differ only in case, the behavior depends on whether the application running on the client accesses the global namespace using Win32 APIs or POSIX APIs.

Integration into Windows:

SAN File System differs from NTFS in its degree of integration into the Windows administrative environment. The differences are as follows:

- Disk Management within the Microsoft Management Console shows SAN File System disks as unallocated.
- SAN File System does not support the use of standard Windows write signature on its disks.
- The SAN File System global namespace cannot be assigned a reserved drive letter.
- Disks used for the SAN File System global namespace cannot be allowed to sleep or hibernate.

SAN File System also differs from NTFS in its degree of integration into Windows Explorer and the desktop. The differences are as follows:

- Manual refreshes are required when updates to the SAN File System global namespace are initiated on the Metadata server (such as attaching a new fileset).
- The recycle bin is not supported.
- Distributed link tracing is not supported. This is a technique through which shell shortcuts and OLE links continue to work after the target file is renamed or moved. Distributed link tracking can help a user locate the link sources in case the link source is renamed or moved to another folder on the same or different volume on the same PC, or moved to a folder on any PC in the same domain.
- Sparse file API and change journaling are not supported. This means that SAN File System provides inefficient support for the indexing service accessible through the Windows ″Search for files or folders″ function.

Privileged clients:

A privileged client is a client machine on which administrative users are given the same privileges on the global namespace as they have on other file systems

available on their system. An administrative user logged in to a privileged Windows client is granted full control over directories, files, and other file system objects created by Windows clients.

If an administrative user logs in to a client that is not a privileged client, that user's privileges for the global namespace are reduced to those of "Everyone."

File names:

File names created on UNIX-based clients using characters that are not valid for the Windows file systems (such as colons, slashes, back slashes, asterisks, question marks, double quote mark, less than, greater than, and pipe) are transformed into valid short names. Applications can use the short name to gain access to files.

Commands:

There is a set of commands that is used to set up, start, and stop the client and to migrate data. These commands are accessible from each client machine. The client commands are separate from the administrative command line interface. The client commands allow users to perform operations on the client systems, while the administrative command line interface allows administrators to administer all aspects of the SAN File System server.

Constraints and limitations:

The following features of NTFS are ***not*** currently supported by SAN File System:

- File compression (for example, a DoubleSpace volume; no support exists for compression on individual files or all files within a folder)
- Encrypted files and directories
- Quotas
- Reparse points (for example, a specialized hook used for certain implementations of hierarchical storage or NTFS junctions)
- Defragmentation and error-checking tools
- Alternate data streams (a concept similar to resource forks in Macintosh)
- Assigning an ACL for the entire drive
- Change journal for file activity
- Scanning all files/directories owned by a particular DID (FSCTL_FIND_FILES_BY_SID)
- Security auditing or SACLs
- Sparse files (as defined by Windows; SAN File System uses sparse allocation techniques, but does not support certain file system APIs required by Windows)

In addition, note these differences:

- Programs that open files using either the 64–bit file ID or the 128–bit object ID (the ″FILE_OPEN_BY_FILE_ID″ option) will fail. This applies to the NFS server bundled with Microsoft Services for UNIX.
- Symbolic links created on AIX clients are handled specially by SAN File System on Windows clients; they appear as regular files whose contents are not accessible and cannot be deleted.
- SAN File System clients acting as CIFS servers do not support opportunistic locks (oplocks), which is an option used by CIFS clients to improve performance through the use of local cache.

Related topics:

- "Clients" on page 10

- "Components" on page 3
- "UNIX-based clients" on page 12

Master console

The *master console* is a serviceability node for SAN File System and other IBM TotalStorage products. SAN File System includes key features that facilitate integration with the master console. In particular, the master console provides key infrastructure for the remote access and service alert features.

The master console for SAN File System is an IBM @server xSeries 305 Type 8673 Model RA1 (and is ordered as a feature of the SAN File System). It uses the following software:

- Microsoft Windows 2000 Advanced Server edition
- IBM Director Server, version 4.1
- IBM Tivoli® Bonus Pack for SAN Management
- Adobe Acrobat, version 5.0
- The PuTTY openssh package

From the master console, you can access the following components:

- SAN File System console through a Web browser.
- Administrative command-line interface through a Secure Shell (SSH) session.
- Any of the engines in the SAN File System cluster through an SSH session.
- The RSA II card for any of the engines in the SAN File System cluster through a Web browser. In addition, you can use the RSA II Web interface to establish a remote console to the engine, allowing you to view the engine desktop from the master console.
- Any of the SAN File System clients through an SSH session, a telnet session, or a remote display emulation package (such as VNC), depending on the configuration of the client.

Typically, you will use the master console to access the engines as well as the SAN File System console or the Administrative command-line interface. However, if necessary, you can attach a keyboard, monitor, and mouse to an engine.

Using the Remote Access feature of the SAN File System, you can initiate a VPN connection to allow a support engineer to remotely access the master console. You can monitor that access and disconnect the session at any time.

Related topics:

- "Service Alert"
- "Remote access" on page 17

Service Alert

Service Alert is a feature that allows IBM service personnel to respond quickly to any problems that occur in your SAN File System server cluster.

When one of the Metadata servers in your server cluster encounters an event that requires a service alert action, it sends a Simple Network Management Protocol (SNMP) datagram to IBM Director, which then parses and converts the trap into a Simple Mail Transfer Protocol (SMTP) e-mail message. The e-mail message is then sent to the customer's SMTP mail server. Finally, the e-mail message is forwarded to IBM personnel for processing.

Figure 2 shows the Service Alert architecture:

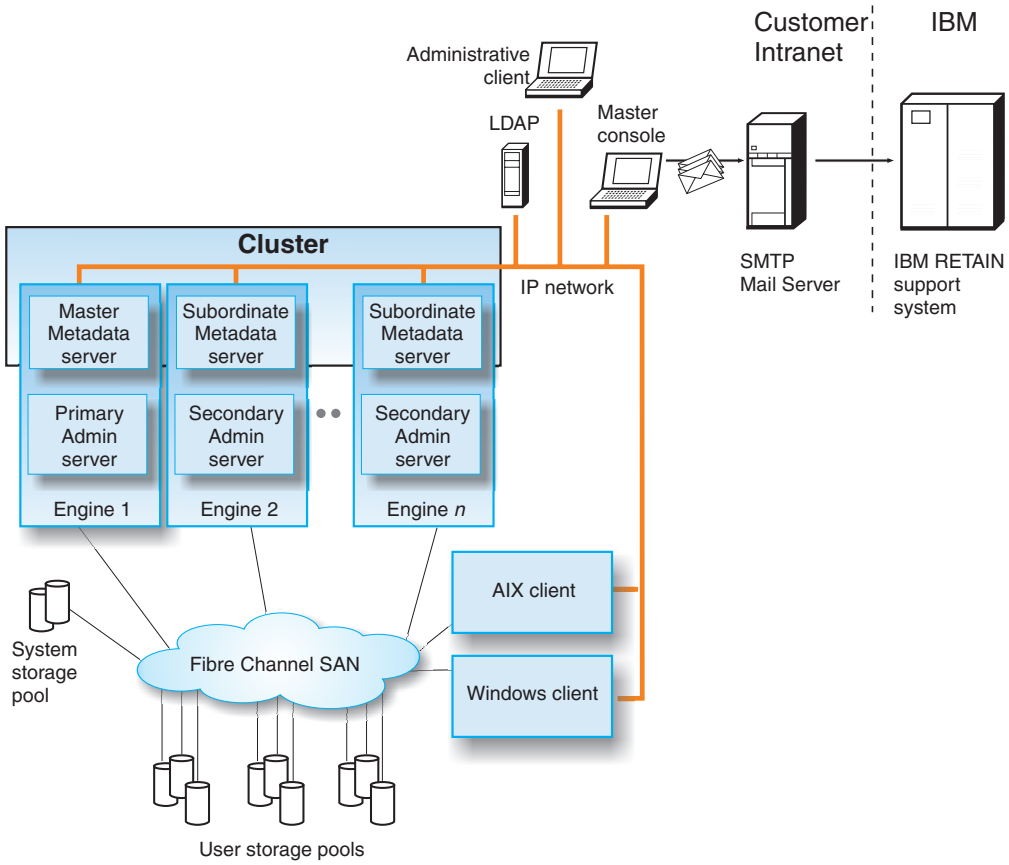

Figure 2. Service Alert architecture

Related topics:

- "Master console" on page 16
- "SNMP" on page 18

Remote access

Remote Access (RA) support is the ability for IBM support personnel who are not located on your premises to assist an administrator or local field personnel in diagnosing and repairing failures with SAN File System. Remote Access support can help to greatly reduce service costs and shorten repair times, which in turn will reduce the impact of any SAN File System failures on your business.

RA provides support personnel with full access to the SAN File System console. This allows the support personnel to query and control the SAN File System Metadata servers and possibly clients, and to access metadata, log, dump, and configuration data.

Remote access does not allow full access with no authentication. The customer must initiate a Virtual Private Network (VPN) connection to allow a support representative to remotely access the master console. From the master console, the support representative can establish a connection to the SAN File System Metadata servers. However, the customer can monitor that access and disconnect the session as desired.

Related topics:

- "Service Alert" on page 16
- "Master console" on page 16
- Chapter 3, "Accessing SAN File System components", on page 19

SNMP

Simple Network Management Protocol (SNMP) is typically used to monitor the health and performance of software, hardware, and networks. SNMP consists of two main components:

- SNMP agents, which are software components that reside on managed devices and collect management information (using Management Information Bases or MIBs). SNMP agents issue traps when SNMP events occur. These traps are sent through User Datagram Protocol (UDP) to an SNMP Manager.
- An SNMP manager, which is a management application (client) that monitors and controls devices using SNMP protocol.

In SAN File System, the Metadata server generates SNMP traps in response to certain events. Note that no SNMP traps are issued from the operating system, hardware, or the Administrative agent.

Note: The RSA II cards can be set up to generate hardware traps as well.

SAN File System administrators can choose to configure which severity levels of events (informational, warning, error, or severe) generate SNMP traps and can specify the SNMP Managers that receive the traps. When an event occurs with a severity level that causes an SNMP trap, SAN File System sends the trap and logs the event in the cluster log.

Note: SAN File System supports asynchronous monitoring through traps but does not support SNMP GETs or PUTs for active management. The SNMP Manager cannot manage SAN File System.

Not all events in SAN File System generate traps. Examples of events that might generate SNMP trap messages include the following:

- When a server executes a change in state
- When a server detects that another server is not active
- When the size of a fileset reaches a specified percentage of its capacity

Related topics:

- "Service Alert" on page 16

Chapter 3. Accessing SAN File System components

There are two types of access that a support representative can utilize to access the components of the SAN File System:

- Using the master console to access the SAN File System. From the master console, you can access the following components:
  - SAN File System console through a Web browser.
  - Any of the engines in the SAN File System cluster through a Secure Shell (SSH) session. From the SSH session, you can access the Administrative command-line interface
  - The RSA II card for any of the engines in the SAN File System cluster through a Web browser. In addition, you can use the RSA II Web interface to establish a remote console to the engine, allowing you to view the engine desktop from the master console.
  - Any of the SAN File System clients through an SSH session, a telnet session, or a remote display emulation package (such as VNC), depending on the configuration of the client.

  Using the Remote Access feature of the SAN File System, the customer can initiate a VPN connection to allow a service representative to remotely access the master console. The customer can monitor that access and disconnect the session at any time. The Remote Access feature allows a service representative to access the master console through one (or both) of the following methods:
  - Establishing an SSH session with the master console over the VPN. From this session, the service representative can establish another SSH session to access the Administrative command-line interface of any of the engines in the SAN File System cluster, or a SAN File System client (if an SSH package is installed and configured on that client). The service representative can also telnet to a SAN File System client if telnet is enabled on that client.
  - Establishing a remote console with the master console over the VPN. This allows the support representative to remotely view the desktop of the master console and perform all of the tasks as if the service representative were working at the master console itself.

    Note: To establish a remote console, a remote display emulation package, such as Virtual Network Computing (VNC), must be installed and configured on the master console by the customer. A VNC application is available at this Web site:

    www.realvnc.com
- Locally attaching a keyboard, monitor, and mouse (or KVM switch) to any of the hardware components of the SAN File System, including the engines in the cluster or any of the SAN File System clients.

  The engines in the cluster are not shipped with a keyboard, monitor, or mouse. Typically, you will use the master console to access the engines as well as the SAN File System console or the Administrative command-line interface. However, if necessary, you can attach a keyboard, monitor, and mouse to an engine.

Note: Depending on the proximity of the master console to the engines in the SAN File System cluster, you may need to locally attach a customer-supplied keyboard, monitor, and mouse to the engine before attempting to service the engine.

Related topics:

- "Remotely accessing the master console"
- "Accessing the Administrative server through a browser" on page 22
- "Accessing an engine through SSH" on page 23
- "Accessing the RSA II adapter" on page 24
- "Accessing a client through SSH" on page 25
- "Accessing a client through telnet" on page 26
- "Accessing a client through a remote console utility" on page 27

Remotely accessing the master console

Prerequisites:

Before initiating a VPN connection between the customer and IBM representative, the following requirements must be met:

- The master console must have a connection to the Internet.
- A Windows user account for the support representative must be set up on the master console.
- If the support representative has a need to access the SAN File System console or the RSA II Web interface remotely, a remote display emulation package, such as Virtual Networking Computer (VNC) server, must be installed and running on the master console.
- The customer must provide a user ID and password for access.
- A maintenance agreement must be established between the customer and IBM or the product must be under software warranty.

Context:

The master console is used to set up a VPN connection between the customer location and IBM support representatives. The customer initiates the connection and has the ability to monitor and control the connection.

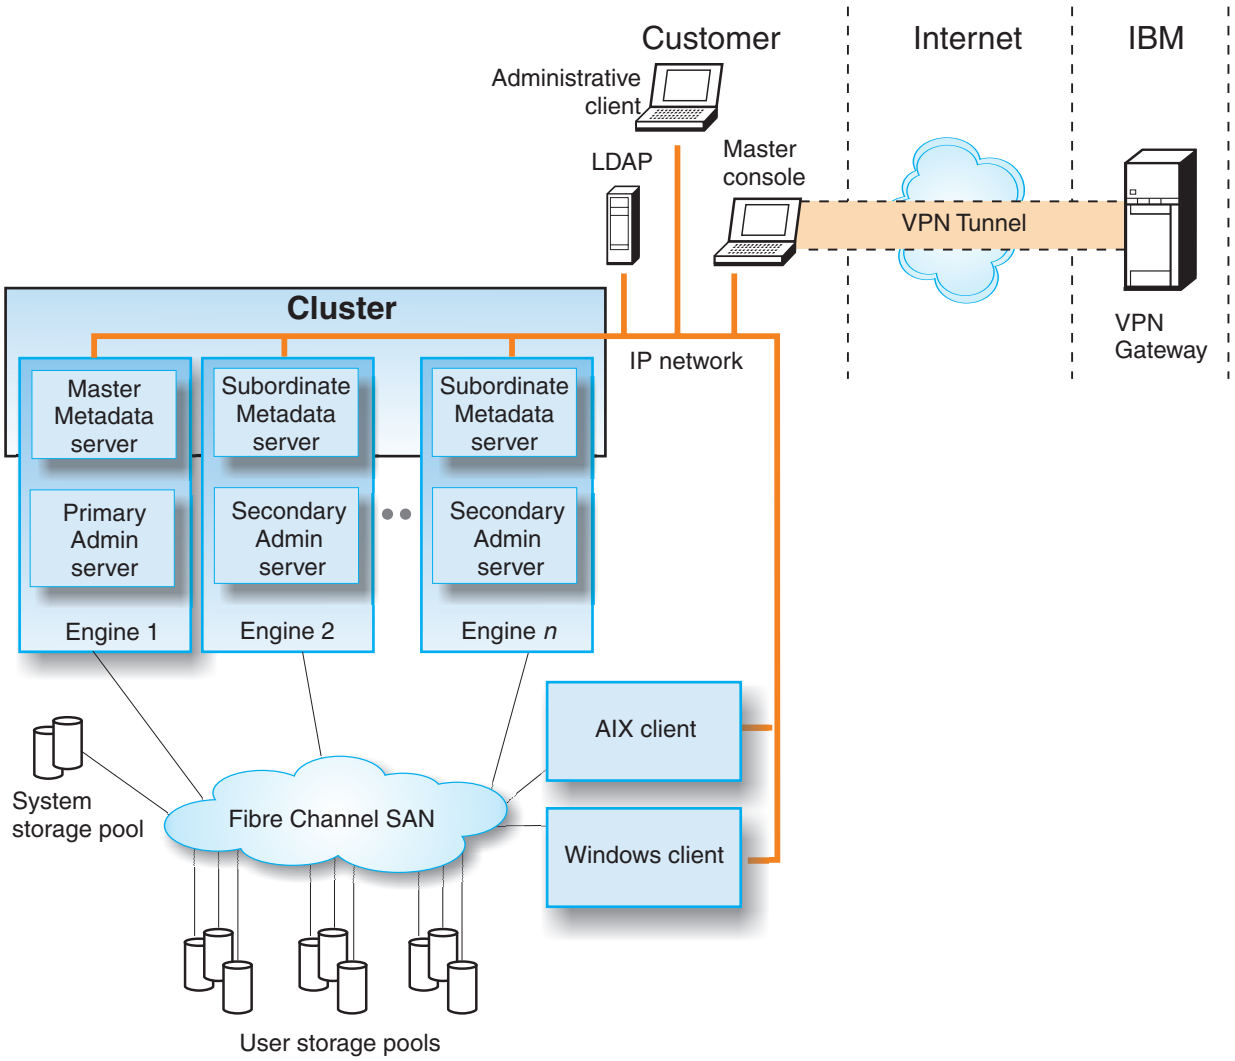

Figure 3. VPN connection between the customer and IBM support representative

Steps:

The following steps describe how a support representative connects to the master console remotely through a VPN connection.

1. The customer logs into the master console. The customer can access the master console directly (using the keyboard, monitor, and mouse) or remotely though another computer on the same LAN.
2. The customer establishes a secure VPN connection from the master console through the VPN gateway to a previously designated VPN server within the IBM intranet. The customer establishes the connection using the IBM connection manager and obtains a connection ID. The IBM connection manager icon is located on the master console desktop.
3. The customer provides the connection ID to the service representative. Each time the customer starts a VPN session, a unique connection ID is created.
4. The service representative connects to the previously designated VPN server within IBM using either a telnet client or a secure shell (SSH) client, such as PuTTY. The service representative uses the connection ID provided by the customer to access the active VPN tunnel.

5. The service representative connects to an account on the master console over the VPN connection to access to the master console. The service representative then establishes a second connection to the VPN server. If a remote display emulation package is available, the service representative can use this connection to establish a remote console to the master console.

Result:

This connection provides a service representative with the ability to log on to these devices or interfaces:

- Each of the engines in the SAN File System cluster. The support representative can query and control the engines at the operating system level by initiating an SSH session with the engine. This requires that a Linux user account be set up on each of the engines in the cluster.
- The Administrative command-line interface. The support representative can query and control the SAN File System Metadata servers, and access metadata, log, dump, and configuration data. This requires that a SAN File System administrative user account be set up for the support representative.
- SAN File System clients. The support representative can query and control clients at the operating system level by initiating either an SSH session or a telnet session with the client (if an SSH or telnet application is installed and running on the client). This requires that an operating system user account be set up on each of the clients to which the support representative will need access.
- SAN File System console and RSA II Web interface (if a remote display emulation package is installed and running).

The customer can monitor all activity performed by the service representative. The customer can either run a remote desktop package from another machine to observe the master console desktop, view the master console SSH log file to see the results of all activity, or watch directly from the monitor on the master console. In addition, the customer has the option to disconnect the VPN session at any time.

Related topics:

- "Master console" on page 16

Accessing the Administrative server through a browser

Prerequisites:

Before accessing the Administrative server through a browser from the master console, the following requirements must be met:

Note: You can access the SAN File System console from any computer that is on the same LAN as the SAN File System cluster.

- A SAN File System administrative user account must be set up for use in signing on to the SAN File System console.
- If the service representative is accessing the master console remotely:
  - The customer must have previously initiated a VPN connection with the service representative.
  - The service representative must have established the VPN connection and used a remote display emulation package, such as VNC, to remotely view the desktop of the master console.

Steps:

Perform these steps to access the SAN File System console through a Web browser.

1. From the master console, open a Web browser and type the URL for the primary Administrative server

   https://*primary_administrative_server***:7979/tank/console**

   where *primary_adminstrative_server* is the host name or IP address of the engine hosting the primary Administrative server (as well as the master Metadata server).

   Note: If you enter a location in the Web browser for an Administrative server other than the primary Administrative server, the request is redirected to the primary Administrative server. If the master Administrative server is not available, the console for the secondary Administrative server is displayed. However, some commands require that the primary Administrative server be available; these commands will not complete successfully.
2. From the SAN File System console welcome page, enter a SAN File System administrative user name and password to sign on.

Result:

From the SAN File System console, you can manage and view information about engines, Metadata servers, and clients. See the *Administrator's Guide and Reference* for more information about the tasks that can be performed from the SAN File System console.

Related topics:

- "Administrative server" on page 8
- "Remotely accessing the master console" on page 20

Accessing an engine through SSH

Prerequisites:

Before accessing an engine using SSH, the following requirements must be met:

- A SAN File System administrative user account must be set up for use in signing on to the SAN File System console.
- A Linux account must be set up on the engine to be accessed for use in signing on to the SSH session.
- If the service representative is accessing the master console remotely:
  - The customer must have previously initiated a VPN connection with the service representative.
  - The service representative must have established the VPN connection.

Steps:

Perform these steps to establish an SSH session to any of the engines in the SAN File System cluster.

1. From the master console, use one of these methods to access the engine:
   - Open a shell prompt and type **putty.exe -ssh** *engine_IP_address*, where *engine_IP_address* is the IP address of the engine to be accessed.

Note: If you used SSH to establish a remote session with the master console, type this command from that session to establish an SSH session between the master console and the engine.

- Click **Start→Programs→PuTTY→PuTTY**.
  a. Type the IP address of the engine to be accessed.
  b. Select SSH as the protocol.
  c. Click **Open**.

2. After the session is established, log in using a Linux user ID and password.

Result:

After connecting to the engine, you can perform these activities:

- Access the SAN File System administrative command-line interface (CLI) to run SAN File System commands. These commands provide the ability to manage engines, Metadata servers, and Administrative servers.
- Access operating-system commands to enable or disable tracing, obtain dumps, and stop or start applications.

Related topics:

- "Remotely accessing the master console" on page 20
- "Administrative commands" on page 158

Accessing the RSA II adapter

Prerequisites:

Before accessing the RSA adapter for an engine, the following requirements must be met:

- An RSA user account must be set up for use in accessing the RSA II Web interface on the engine.
- A SAN File System administrative user account must be set up for use in signing on to the SAN File System console.
- If the service representative is accessing the master console remotely, the following conditions must be met:
  – The customer must have previously initiated a VPN connection with the service representative.
  – The service representative must have established the VPN connection and used a remote display emulation package, such as VNC, to remotely view the desktop of the master console.

Steps:

Perform these steps to access the RSA II Web interface for an engine:

1. Open a browser from the remote console on the master console and type the URL for an RSA adapter:

   http://*RSA_II_web_address***/**where the *RSA_II_web_address* is the IP address of the RSA II card.
2. Log on to the RSA II interface using a valid RSA user name and password.

Result:

From the left navigation pane, you can choose to perform tasks such as:

- Stopping and restarting the engine
- Viewing vital product data for the engine
- Accessing the BIOS and firmware for the engine
- Updating the firmware for the RSA card
- Accessing the RSA adapters of other engines in the SAN File System cluster

Note: If, when you view the System Health Summary page, all components are listed as unavailable, make sure the PCI-riser card assembly is firmly seated within the engine.

You can also start a remote video console that can be used to redirect the engine console to the master console. In addition, you can use the remote control functionality to assign the CD-ROM or diskette drive from the master console to be used by the engine.

Note: To use the remote control functionality of the RSA adapter, you must be signed on with an ID defined in the RSA II card that has read/write access to the RSA II adapter.

For more information about using the RSA card, see the *Remote Supervisory Adapter II User's Guide*, which is available from this Web site (search for Remote Supervisor Adapter II from the Search Technical Support link):

www.ibm.com/storage/support/

Related topics:

- "Remotely accessing the master console" on page 20

Accessing a client through SSH

Prerequisites:

Before accessing a client using SSH, the following requirements must be met:

- An operating system account must be set up for use in signing on to the client to be accessed.
- An SSH software package, such as PuTTY, must be installed and running on the client to be accessed.
- If the service representative is accessing the master console remotely, the customer must have previously initiated a VPN connection with the service representative.

Steps:

Perform these steps to establish an SSH session with any of the SAN File System clients.

1. From the master console, use one of these methods to access the client:
   - Open a shell prompt and type **putty.exe -ssh** *client_IP_address***,** where *client_IP_address* is the IP address of the client to be accessed.

     Note: If you used SSH to establish a remote session with the master console, type this command from that session to establish an SSH session between the master console and the client.

- Click **Start→Programs→PuTTY→PuTTY**.
  a. Type the IP address of the client to be accessed.
  b. Select SSH as the protocol.
  c. Click **Open**.

2. After the session is established, you can log in using an operating system user ID and password.

Result:

After the connection is established, you can perform these activities:

- If you are accessing an AIX client, you can run SAN File System commands that provide the ability to stop and start the client as well as list client status. If you are accessing a Windows client, you can run the **migratedata** command.
- Access operating-system commands to enable or disable tracing, obtain dumps, and stop or start applications.

Related topics:

- "Client commands" on page 297
- "Remotely accessing the master console" on page 20
- "Service commands and utilities" on page 319

Accessing a client through telnet

Prerequisites:

Before accessing a client using telnet, the following requirements must be met:

- An operating system account for the support representative must be set up on the client to be accessed.
- If the service representative is accessing the master console remotely, the customer must have previously initiated a VPN connection with the service representative.

Steps:

Perform these steps to establish a telnet session with any of the SAN File System clients.

1. From the master console, open a shell prompt and type **telnet** *host*, where *host* is the IP address of the client to be accessed.
2. After the session is established, you can log in using an operating system user ID and password.

Result:

After the connection is established, you can perform these activities:

- If you are accessing an AIX client, you can run SAN File System commands that provide the ability to stop and start the client as well as list client status. If you are accessing a Windows client, you can run the **migratedata** command.
- Access operating-system commands to enable or disable tracing, obtain dumps, and stop or start applications.

Related topics:

- "Client commands" on page 297
- "Remotely accessing the master console" on page 20
- "Service commands and utilities" on page 319

Accessing a client through a remote console utility

Prerequisites:

Before accessing a client using a remote console utility, such as Virtual Network Computing (VNC), the following requirements must be met:

- An operating system account must be set up on the client to be accessed.
- A VNC server must be installed and running on the client to be accessed. In addition, a session password must exist.
- A VNC client must be installed and running on the master console.
- If the service representative is accessing the master console remotely:
  - The customer must have previously initiated a VPN connection with the service representative.
  - The service representative must have established the VPN connection and used a remote display emulation package, such as VNC, to remotely view the desktop of the master console.

Steps:

Perform these steps to establish a VNC session with any of the SAN File System clients.

1. Double-click the VNC viewer icon on the master console.
2. Enter the IP address or host name of the client to be accessed and click **OK**.
3. Enter the session password and click **OK**.

Result:

After the connection is established, you will see the remote desktop of the client. From the VNC connection, you can perform tasks and run commands as if you were physically at the client machine.

Related topics:

- "Client commands" on page 297
- "Remotely accessing the master console" on page 20

Chapter 4. Diagnostic tools

You have several tools available to you when you attempt to diagnose problems with SAN File System components:

- Engine. You can use the power-on self-test (POST), light-path diagnostics, and hardware diagnostic tests to analyze and isolate problems with an engine. In addition, you can view logs to understand where problems might be occurring.

  You can use these tools by locally accessing an engine, or you can use the RSA II Web interface to access the engine remotely.
- Server. The SAN File System provides logs that you can use to view information about the Metadata server and the Administrative server. You can access these logs directly from the engine hosting a server or view a consolidated version of the logs through either the administrative command-line interface or SAN File System console.

  You can also use SAN File System tools as well as Linux operating system tools to generate and analyze traces and dumps.
- Client. From a SAN File System client, you can use SAN File System commands and operating system commands to view, capture, and analyze log and trace data, as well as system dumps.

Related topics:

- "Engine diagnostic tools" on page 29
- "Server diagnostic tools" on page 36
- "Client diagnostic tools" on page 42

Engine diagnostic tools

You can use the following diagnostic tools to resolve problems that may occur with an engine:

- Power-on self-test (POST). When you power on an engine, it performs a POST to verify the operation of engine components as well as some of the options installed in the engine.
- Light path diagnostics. If the system-error LED on the front of the engine is lit, you can open the cover and use the Light Path Diagnostics panel located inside the server to isolate problems.
- Hardware diagnostic tests. When you power on an engine, you can press **F2** to display the Diagnostics Programs panel. From this panel, you can run tests to verify that engine components are operating properly.
- Error Logs. There are several error logs that you can view to assist in determining what is causing an error.

Related topics:

- "POST" on page 30
- "Light path diagnostics" on page 30
- "Diagnostic tests" on page 31
- "Engine error logs" on page 32

POST

When you power on an engine, it performs a series of tests to verify the operation of engine components as well as some of the options installed in the engine. This series of tests is called the power-on self-test, or POST.

If POST completes without detecting any problems, a single beep sounds, and the boot panel appears. If POST detects problems, more than one beep sounds and an error message (POST error code) appears on the screen.

Notes:

1. If you have a power-on password or administrator password set, you must type the password and press **Enter**, when prompted, before POST will continue.
2. A single problem might cause several error messages. Therefore, always work to correct the cause of the first error message. Then attempt to complete POST again.

The system error log contains all messages issued during POST and all system status messages from the service processor. In addition, the POST error log contains the three most recent error codes and messages generated during POST.

Related topics:

- "Beep symptoms" on page 377
- "POST error codes" on page 389
- "Engine error logs" on page 32

Light path diagnostics

If the system-error LED on the front of the engine is lit, one or more LEDs on the Light Path Diagnostics panel might also be lit. The Light Path Diagnostics panel is located inside the engine, and you can access the panel by removing the cover. These LEDs remain lit even when the engine shuts down, as long as the power supply is operating properly.

Diagnostics panel LEDs:

The following illustration shows the LEDs on the Light Path Diagnostics panel.

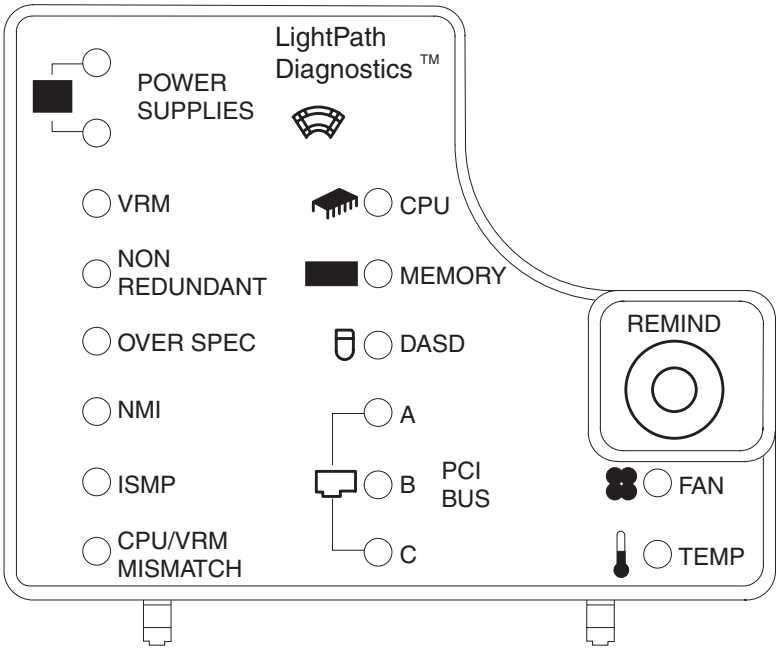

Figure 4. Diagnostics panel LEDs

Remind button:

You can use the Remind button to place the front panel system-error LED into Remind mode, to indicate that you intend to delay resolution of the error until a later time. In Remind mode, the system-error LED flashes every two seconds. The system-error LED remains in Remind mode until one of the following situations occurs:

- You resolve all known problems.
- You restart the engine.
- A new problem occurs.

Related topics:

- "Light path diagnostic panel LEDs" on page 387

Diagnostic tests

Engine diagnostic programs are stored in read-only memory (ROM) on the system board. You can run these tests to verify that the major components of the engine are operating properly. Error messages generated by the tests are stored in a summarized test log as well as in the test log for the specific test.

Message format:

Messages generated by the diagnostic tests are in the following format:

fff-ttt-iii-date-cc-text message

where:

fff is the three-digit function code that indicates the function being tested when the error occurred. For example, function code 089 is for the microprocessor.

ttt is the three-digit failure code that indicates the exact test failure that occurred.

iii is the three-digit device ID.

date is the date that the diagnostic test was run and the error recorded.

cc is the check value that is used to verify the validity of the information.

text message
is the diagnostic message that indicates the reason for the problem. The format of the text message is:

Function name: *result* (*test-specific string*)

where:

Function name
is the name of the function being tested when the error occurred. This name corresponds to the function code (*fff*) shown in the error code format.

result is the result of the test. It can be one of the following values:

- Passed. The diagnostic test completed without any errors.
- Failed. The diagnostic test discovered an error.
- User Aborted. You stopped the test before it was completed.
- Not Applicable. You specified a diagnostic test for a device that is not present in the engine.
- Aborted. The test cannot proceed. This result can occur, for example, if the system configuration is not correct.
- Warning. A potential problem was discovered during the test. This result can occur, for example, when a device driver is not found.

test-specific string
provides additional information that you can use to analyze the error that was found.

Related topics:

- "Running engine diagnostics" on page 104
- "Diagnostic error codes" on page 381

Engine error logs

The engine provides the following error logs:

- POST error log

  The POST error log contains the three most recent error codes and messages generated during the power-on self-test that occurs when you power on the engine. These messages also appear on the screen during power on.
- System error log

  The system error log contains all of the error messages generated by the engine components. It includes all messages written to the POST error log.
- Test log

  The test log contains error messages generated from the diagnostic programs.

You can view these logs from either the diagnostic programs or from the Configuration and Setup utility.

In addition to these logs, the Remote Supervisor Adapter also provides an error log called the event log. From the RSA II Web interface, you can access the event log. It contains all of the messages in the system log. It also contains information about remote-access and dial-out attempts.

Related topics:

- "Diagnostic tests" on page 31
- "Running engine diagnostics" on page 104
- "Starting the Configuration/Setup Utility" on page 106

OS Present diagnostics

In addition to the firmware diagnostics provided with the engine, the SAN File System provides an additional set of OS Present diagnostics that enable you to test the following adapter cards that are installed in the engine:

- Fibre-channel adapter
- Ethernet adapter
- Remote Supervisor Adapter (RSA) II
- System board Ethernet ports

The OS Present diagnostics check that device drivers are installed and operational. In addition, you can use OS Present diagnostics to view the error log from adapter diagnostics tests; obtain vital product data for adapters, processors, memory, and system board; view information from the RSA II event log; and view a listing of changes that have been made to hardware devices, their firmware, or device driver versions.

You can access the OS Present diagnostics using the **diags** command, which provides a menu-driven interface to the list of devices in the engine. For each available device, you can run diagnostic tests, as well as view information about the device and the tests that are run.

In addition, OS Present diagnostics provide tools for viewing logs directly from the command line:

- The **delog** command displays the latest diagnostics error log.
- The **elog** command displays the RSA II log.
- The **vpdlog** command displays vital product information about the devices.

Related topics:

- "The diags command"
- "The delog command" on page 34
- "The elog command" on page 34
- "The vpdlog command" on page 35

The diags command

You can use the **diags** command to display information about all devices that are installed in the engine. In addition, you can also run diagnostic tests for each available adapter card in the engine. The list of tasks available for the adapter card is based on the operational mode of the engine:

- In normal operations mode, you can perform these tasks for a device:
  - Flash the blue "identity" LED on this engine for 2 minutes. This allows you to quickly find the engine in a rack.

  - View the diagnostic error log for this engine. The diagnostic error log displays the log from the last time that the diagnostic tests were run.
  - View the Remote Supervisor Adapter II event log.
  - View the service history log. The service history log provides a list of all changes that have been made to the devices in the engine, such as when a device was replaced or its firmware was updated. The log is located in /var/log/jupiter/service-history.txt and is cumulative. It continues to be updated until it is deleted or the hard drive fails.
  - Reboot or shut down the engine. You will need to reboot the engine when switching between normal operations mode and service mode. You can also reboot the engine to ensure that the latest device information is available.
- In service mode, you can perform all of the tasks that are available in normal operations mode. In addition, you can perform these tasks:

  Note: You must be logged in as root to run the diagnostic tests.
  - Run the adapter test.
  - Loop the adapter test until an error occurs or until you manually stop the test.
  - Stop a looping test.

Related topics:

- "Running OS Present diagnostics" on page 105
- "The delog command"
- "The elog command"
- "The vpdlog command" on page 35

The delog command

The **delog** command lists the latest diagnostics error log. This log contains any errors detected when you run tests with the **diags** command.

By default, the **delog** command writes the log to /tmp/diag-errorlog.txt. However, you can also run **delog stdout** to display the log on the console or **delog** *filename* to write the log to a specified file.

Related topics:

- "The diags command" on page 33
- "The elog command"
- "The vpdlog command" on page 35

The elog command

The **elog** command lists the RSA II log. By default, the **elog** command writes the log to /tmp/eventlog.txt. However, you can run **elog stdout** to display the log on the console or **elog** *filename* to write the log to a specified file.

Related topics:

- "The diags command" on page 33
- "The delog command"
- "The vpdlog command" on page 35

The vpdlog command

The **vpdlog** command lists all vital product data that can be collected for the adapter cards in the engine. For each adapter, the **vpdlog** command displays information such as:

- Name
- Model number
- Size or capacity of the device
- IBM serial number
- IBM field-replaceable unit (FRU) number
- IBM part number
- Engineering change level
- Firmware code levels (1 to 3)
- Network address (MAC address or World Wide Name, as appropriate)

By default, the **vpdlog** command writes the log to /tmp/vpdlog.txt. However, you can also run **vpdlog stdout** to display the log on the console or **vpdlog** *filename* to write the log to a specified file.

Related topics:

- "The diags command" on page 33
- "The delog command" on page 34
- "The elog command" on page 34

RSA II Web interface

A Remote Supervisor Adapter (RSA) II is installed in each engine in the cluster. Each RSA II is connected to the IP network. In addition, the RSA II in each engine in the cluster creates a daisy chain through a private Advanced System Management (ASM) interconnect network. For more information about the RSA II, refer to the *Remote Supervisor Adapter II User's Guide*, which is available on the IBM Web site.

Using the RSA II Web interface, you can perform the following functions:

- Access system status, which allows you to view the health of the engine, the image of the last operating-system-failure screen capture, and a list of users who are logged in to the RSA II.
- View the event logs, which contain all entries from the engine system error log. It also contains information about remote-access attempts and dial-out events.
- View vital product data for the engine, including BIOS information and engine component vital product data.
- Remotely power on or restart the engine.
- Redirect the engine video console to your computer. Using the remote control function, you can also use your computer disk drive or disk image as a drive on the engine.
- Change the engine boot sequence for the next restart or attempt a Preboot Execution Environment (PXE)/Dynamic Host Configuration Protocol (DHCP) network startup.
- Update the firmware on the RSA II.
- Access other system-management processors and adapters on the ASM interconnect network. This includes RSA II cards, ASM processors, ASM PCI

adapters, and integrated system management processors (ISMPs). Collectively, these system-management processors and adapters are called *service processors*.

- View and change the RSA II system settings and RSA II clock.
- Configure RSA II login profiles.
- Configure remote alerts and remote alert recipients and establish the events that will be considered to be alerts.
- Configure RSA II serial port settings.
- Configure RSA II network interfaces. This includes the settings for the Ethernet connection on the RSA II, which enables remote access over a Web browser, network protocols used by the RSA II, and secure socket layer (SSL) configuration.
- Back up and restore the RSA II configuration. In addition, you can restore RSA II defaults.
- Restart the RSA II.

Related topics:

- "Logging on to the RSA II" on page 106

Server diagnostic tools

You can use the following tools to diagnose problems that you are having with either the Metadata server or the Administrative server:

- Logs. The SAN File System provides several logs that you can use to view information about Metadata servers and Administrative servers.
- Tracing. The SAN File System provides the ability to trace both the Metadata server and the Administrative server to diagnose problems.
- Dumps. You can use dump capabilities provided by the SAN File System, as well as the dump capabilities of the Linux operating system, to diagnose problems.

Related topics:

- "Logs" on page 36
- "Metadata server tracing" on page 40
- "SAN File System dump capability" on page 41
- "Linux dump capability" on page 41

Logs

The following logs for the Metadata server and the Administrative server are stored on the engine hosting those servers.

Table 1. SAN File System message log files

Log	File name	Location	Maximum file size
Administrative log	cimom.log	/usr/tank/admin/log	–
Audit log	log.audit	/usr/tank/server/log	250 MB
Console log	console.log	/usr/tank/admin/log	–
Dump log	log.dmp	/usr/tank/server/log	–
Failover log	log.failover	/usr/tank/server/log	–

Table 1. SAN File System message log files (continued)

Security log	security.log	/usr/tank/admin/log	–
Server log	log.std	/usr/tank/server/log	250 MB
Standard error	stderr.log	/usr/tank/admin/log	–
Standard out	stdout.log	/usr/tank/admin/log	–
Trace log	log.trace	/usr/tank/server/log	250 MB

Notes:

1. Although log.audit, log.std, and log.trace have a maximum file size of 250 MB, the SAN File System actually stores 500 MB of data for each of these logs. When either of these logs reaches its maximum size, it is renamed to include the .old extension. If a file by that name already exists, the existing file is overwritten. Then the log is cleared so that it can start accepting new messages again.
2. The log.dmp file starts over for either of these occurrences:
   - The start of each day
   - The file reaches a size of 1 MB.

When you display these logs from the master Metadata server using either the Administrative command-line interface or the SAN File System console, you actually see a consolidated view of all of the logs from each engine in the cluster.

Notes:

1. The consolidated view of the server message log is called the cluster log.
2. You can also display the event log. This log is actually a subset of the messages stored in the cluster log. It contains only those messages that have a message type of event.

Related topics:

- "catlog" on page 169
- "Administrative log"
- "Audit log" on page 38
- "Security log" on page 39
- "Server log" on page 40

Administrative log

The administrative log contains messages generated by the Administrative server. If you display the administrative log from either the Administrative command-line interface or the SAN File System console, all administrative logs on all engines in the cluster are consolidated into a single view.

Fields:

Log entries contain the following fields:

Message ID
A unique identifier for the message.

Severity level
Indicates whether the entry is an informational, warning, error, or severe message.

Message type
Specifies whether the message is a normal log entry or one that was generated as a result of an event on the Administrative server.

Administrative server ID
A unique identifier for the Administrative server on which the command was issued.

Timestamp
A date followed by a local time.

Message
A textual explanation of the message.

Example:

The following example illustrates the consolidated view of the administrative log that is displayed through the Administrative command-line interface:

```
CIMServer: Info   Normal mds_engine_0 May 16, 2003 7:27:33 AM
  Namespace \root\cimv2 initialized
CMMOM0411I Info   Normal mds_engine_0 May 16, 2003 7:27:33 AM
  Authorization is not active
CMMOM0901I Info   Normal mds_engine_0 May 16, 2003 7:27:33 AM
  IndicationProcessor started
CMMOM0906I Info   Normal mds_engine_0 May 16, 2003 7:27:33 AM
  No pre-existing indication subscriptions
CMMOM0404I Info   Normal mds_engine_0 May 16, 2003 7:27:33 AM
  Security server starting on port 5989
CMMOM0402I Info   Normal mds_engine_0 May 16, 2003 7:27:33 AM
  Platform is Unix
```

Related topics:

- "catlog" on page 169

Audit log

The Administrative audit log contains all administrative actions issued to the SAN File System. It contains a record of every command issued by a SAN File System administrator, from either the Administrative command-line interface or the SAN File System console, that changes the state of the Metadata server in some way.

The Administrative audit log contains a record of every command issued by a SAN File System administrator, from either the command-line interface or the SAN File System console, that changes the state of the Metadata server in some way.

Each Administrative server has its own audit log. If you display the audit log from either the Administrative command-line interface or the SAN File System console, all audit logs on all engines in the cluster are consolidated into a single view.

Fields:

Log entries contain the following fields:

Timestamp
A date followed by a local time.

Severity level
Set to a value of Informational (console) if the command succeeded. Otherwise, it is set to a value of Error.

Message ID
A unique identifier for the message.

Message type
Set to Audit. This field is contained in the audit log, but it is not displayed in the consolidated view.

Metadata server ID
A unique identifier for the Metadata server on which the command was issued.

Message
Contains the user name of the SAN File System administrator who issued the command followed by a functional replica of the message itself.

Example:

The following example shows a message from the consolidated view of the audit log that is displayed through the Administrative command-line interface:

```
2003-04-21 18:36:32  INFORMATIONAL  HSTAD0083I  A  mds_engine_0
User Name: jozvold  Command Name: MasterServiceAddServer
Parameters: SYSTEMCREATIONCLASSNAME=STC_Cluster
SYSTEMNAME=testnode CREATIONCLASSNAME=STC_MasterService
NAME=MasterService CLUSTERPORT=10001 IP=9.29.25.136
Command Succeeded.
```

Related topics:

- "catlog" on page 169

Security log

The security log displays the administrative user login activity for the Administrative server. If you display the security log from either the Administrative command-line interface or the SAN File System console, all security logs on all engines in the cluster are consolidated into a single view.

Fields:

Log entries contain the following fields:

Message ID
A unique identifier for the message.

Severity level
Indicates whether the entry is an informational, warning, or error message.

Administrative server ID
A unique identifier for the Administrative server on which the command was issued.

Message
A textual explanation of the message.

Example:

The following example illustrates the administrative log displayed through the Administrative command-line interface:

```
CMMOM0302I Info  mds_engine_0 May 19, 2003 9:21:17 AM
     User respey on client {1} could not be authenticated
CMMOM0302I Info  mds_engine_0 Jun 13, 2003 1:51:40 PM
     User jkaminski on client {1} could not be authenticated
CMMOM0302I Info  mds_engine_0 Jun 20, 2003 5:41:36 PM
     User fstock on client {1} could not be authenticated
```

Related topics:

- "catlog" on page 169

Server log

The server log contains operational information. Each Metadata server in the SAN File System cluster has its own log. If you display the server log from either the Administrative command-line interface or the SAN File System console, all server logs on all engines in the cluster are consolidated into a single view. The consolidated view of the server message log is called the cluster log.

Fields:

Log entries contain the following fields:

Timestamp
A date followed by a local time.

Severity level
Indicates whether the entry is an informational, warning, error, or severe message.

Message ID
A unique identifier for the message.

Message type
Specifies whether the message is a normal log entry or one that was generated as a result of an event on the Metadata server.

Metadata server ID
A unique identifier for the Metadata server on which the command was issued.

Message
A textual explanation of the message.

Example:

The following example shows a message from the cluster log that is displayed through the Administrative command-line interface:

```
2003-04-16 12:55:50  INFORMATIONAL  HSTPG0009I   N
     msd_engine_0  Using IP 9.38.203.26 port 10192
     for Group Services messages.
```

Related topics:

- "catlog" on page 169

Metadata server tracing

You can enable tracing for the Metadata server through the Administrative command-line interface using the trace command. This command enables you to control:

- When tracing begins and ends.
- The components within the Metadata server for which tracing will occur.
- The level of detail (verbosity) to show during tracing.

Tracing generates a trace log called log.trace, that is stored in /usr/tank/server/log. Like other log files, this file has a maximum file size of 250

MB. When it reaches this size, the SAN File System creates a copy called log.trace.old and clears log.trace. If log.trace.old already exists, it is overwritten.

Note: A minimum level of tracing always occurs for the Metadata server, so this file always exists.

Related topics:

- "legacy" on page 328

SAN File System dump capability

The SAN File System provides three ways that you can gather diagnostic data about the engines in the cluster as well as the Metadata servers that run on those engines:

- Run the **collectdiag** command from the Administrative command-line interface. The collectdiag command allows you to collect information about an engine in the cluster.
- Click **Maintain System→Collect Diagnostics** from the SAN File System console. The Collect Diagnostics task is the console equivalent of running the **collectdiag** command from the Administrative command-line interface.
- Use the One-Button Data Collection Utility.

Related topics:

- "Administrative commands" on page 158
- "One-button data collection utility" on page 45

Linux dump capability

You can use the Linux operating system process dump capabilities to help with debugging and problem determination of the SAN File System. In addition, you may need to force a core dump of the Metadata server process so that you can package the data and send it to other support personnel for review. Follow these steps to force a core dump:

1. Use the **ulimit** shell command to set the size of the allowable core dump file size to be unlimited.

   Note: You can use the **ulimit** shell command with the –a parameter to verify the current allowable limit.

   ```
   ulimit –c unlimited
   ```

2. Use the Linux **kill** command to terminate the Metadata server process:

   kill –6 *PID*

   where *PID* is the process ID for the Metadata server process.

   Note: You will need to terminate all SAN File System processes that are currently running. Typically, using the **kill** command against the parent process will also terminate all child processes.

3. The **kill** command produces a file call core.*PID* in the directory where you entered the **kill** command.

The Metadata server runs in user space. Therefore, a problem with the Metadata server should not crash the Linux kernel. You should not need to analyze kernel dumps on Linux for the Metadata server.

Client diagnostic tools

You can use the logging and tracing capabilities provided by the SAN File System to perform problem determination with SAN File System clients on the AIX operating system and the Windows 2000 operating system. In addition, both operating systems provide the ability generate and analyze system dumps.

Related topics:

- "AIX client logging and tracing" on page 42
- "AIX client dump capability" on page 43
- "Windows 2000 client logging and tracing" on page 44
- "Windows 2000 client dump capability" on page 45

AIX client logging and tracing

Use the **stfsdebug** command and the syslog facility to enable tracing and logging on the SAN File System AIX client.

Syslog facility:

The SAN File System client generates both log and trace messages, which are routed through the syslog facility on the AIX operating system. The syslog facility captures log and trace output from the kernel as well as other operating system services.

By default, the syslog facility discards all kernel output. However, you can configure the syslog facility to specify a destination for the messages by modifying /etc/syslog.conf.

- Specifying a file as the destination.

  You can specify a file to receive kernel messages, such as /var/adm/ras/messages. To specify that file, perform the following steps:

  1. Create /var/adm/ras/messages if it does not already exist. You can use the AIX touch command to create an empty file.
  2. Edit /etc/syslog.conf.
  3. Insert this line:

     ```
     kern.debug /var/adm/ras/messages
     ```
  4. Restart the syslogd daemon.

     ```
     kill -hup syslogd_PID
     ```

     Refer to the *AIX Commands Reference* for more information about the syslogd daemon.

- Specifying the console as the destination.

  Note: If you specify the console as the destination, messages are also written to /var/spool/mqueue/syslog.

  To specify the console as the destination for kernel messages, perform the following steps:

  1. Edit /etc/syslog.conf.
  2. Insert this line:

     ```
     kern.debug /dev/console
     ```
  3. Restart the syslogd daemon.

```
kill -hup syslogd_PID
```

Refer to the *AIX Commands Reference* for more information about the syslogd daemon.

When you specify kern.debug as shown in the previous examples, all levels of kernel output are routed because debug is the lowest priority level of kernel output. You could specify a different level of output, such as kern.info to show just informational messages.

Stfsdebug command:

You can use the **stfsdebug** command to enable tracing for an AIX client. In addition, you can specify which components (called classes) are traced as well as the level of detail to include. You can also use the **stfsdebug** command to query the current status of all trace classes.

The **stfsdebug** command requires the full pathname of the SAN File System kernel module loaded on the client machine, which you can find by viewing the client configuration file.

Trace output enabled by the **stfsdebug** command is sent to the syslog facility.

Example:

The following example messages show the format of log messages:

```
Apr 21 07:43:50 aixclient1 unix: STFS: disk configuration process created
   with PID = 13348
Apr 21 07:43:50 aixclient1 unix: STFS: cleaner process created with PID 12028
Apr 21 07:43:50 aixclient1 unix: STFS: CSM process created with PID 10860
```

The following example messages show the format of trace messages:

```
Apr 28 13:17:09 aixclient1 unix: STFS: 1051550182.439290  50337 STFS
   traceBuf_daemonize: going to sleep till shutdown
Apr 28 13:17:09 aixclient1 unix: STFS: 1051550182.448769 196267 STFS CSM
   OS-dependent services initialized.
Apr 28 13:17:09 aixclient1 unix: STFS: 1051550182.448827 196267 STFS Pager
   Strategy initialized.
Apr 28 13:17:09 aixclient1 unix: STFS: 1051550182.448875 196267 STFS GFS
   hooks initialized.
Apr 28 13:17:09 aixclient1 unix: STFS: 1051550182.448969 196267 STFS
   doInit(): system Initialized
```

Related topics:

- "stfsdebug" on page 345

AIX client dump capability

You may need to use the dump capabilities if the SAN File System client virtual file system (AIX) hangs, meaning that it is still running but will not respond to commands. Perform one of these steps to initiate a kernel dump.

Note: The SAN File System client runs in kernel space.

- If you can access the client machine remotely (using telnet or ssh) to obtain a shell prompt, issue the command **sysdumpstart**. For information about using the **sysdumpstart** command, refer to *AIX Version 5.1 Commands Reference, Volume 5*, which is available from the IBM Web site.

- If you cannot access the client machine remotely, you can initiate a kernel dump using the reset button on the front panel of the machine. See the documentation that was provided with the client machine to determine how to reset it.

By default, the dump file is saved as /var/adm/ras/vmcore.*n*, where *n* is a number that is incremented each time a dump file is created.

Windows 2000 client logging and tracing

A SAN File System Windows 2000 client provides two types of messages:

- Log messages that provide information, warnings, and errors of general interest to administrators and support personnel.

  Log messages are written to the standard system logging interface, the Windows Event Log. In addition to the operating system messages, the Windows Event Log contains messages generated by the SAN File System.

  You can use the Event Viewer to list messages from the Event Log. If you double-click a message from the Event Viewer, you can find more detailed information about that message. You can also use the Event Viewer to filter messages by message type, source of the message, or according to a specified range of time. You can also dump events to a text file, which is useful for sending problem determination data to remote support personnel.

  The following example messages show the format of the log messages:

```
4/21/2003 7:32:03 PM  Stfs  Error        None  9  N/A  WINCLIENT1
   HSTCW0009E: Couldn't contact server at IP address <18.18.18.99:11190>
4/21/2003 7:32:36 PM  Stfs  Information  None  8  N/A  WINCLIENT1
   HSTCW0008I: Contacted server at IP address <18.18.18.99:11190>.
4/21/2003 7:32:02 PM  Stfs  Information  None  1  N/A
   WINCLIENT1  HSTCW0001I: SAN File System client started successfully.
```

- Trace messages that consist of extensive low-level tracing output about client functions and internal data.

  You can use the **stlog** command to enable and control tracing for a client. However, you should use the **stlog** command with care; enabling full tracing can significantly impact the performance of the SAN File System.

  If tracing is enabled, the SAN File System writes trace messages to a file named c:\Program Files\IBM\Storage Tank\client\log\sanfs.log. This file contains tracing output only for the SAN File System client. It does not contain information for the operating system or any other applications.

  Note: By default, minimal tracing is enabled.

  Enabling detailed tracing

  To enable detailed tracing, you must provide a path and filename to a file to use for the detailed trace log:

  1. Start the Windows registry editor.
  2. Navigate to the following registry key:

     \\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Stfs\Trace
  3. In this key there is an empty string value named FileName. Double-click on FileName and change its value in the string editor to a valid path and file name for the trace log (for example, c:\stfs.log)
  4. Click OK.
  5. Reboot for the changes to take effect.

The following example messages show the format of trace messages:

```
#E   8125026652|80BF9DA0 Reassert TStreamSocket::Disconnect:1911
     84860B68 192.168.10.6:10290       CheckStatus failed:
     STATUS_CONNECTION_ACTIVE (C000023B)
     8125026715|8122C9E0 ckground TBackground::Main:1301
     F4A8EF88               Active count: 00000001
     8125026715|8122C9E0 ckground TBackground::Main:1305
     F4A8EF88               Active work items: 00000000
     8125026715|8122C9E0 ckground TBackground::Main:1308
     F4A8EF88               Active delayed work items: 00000000
     8125031715|8122C9E0 ckground TBackground::Main:1301
     F4A8EF88               Active count: 00000001
     8125031715|8122C9E0 ckground TBackground::Main:1305
     F4A8EF88               Active work items: 00000000
     8125031715|8122C9E0 ckground TBackground::Main:1308
     F4A8EF88               Active delayed work items: 00000000
     8115545496|FE2D7020          TSc::Reference
     FE484008 Fil:<23456.3.1342760.0>  ReferenceCount 3, CsmHandle_Held
```

Windows 2000 client dump capability

You can configure the operating system to generate a dump file if the SAN File System client on the Windows operating system terminates abnormally. By default, the file is C:\WINNT\memory.dmp. You can also configure where the file is saved.

In cases where the SAN File System client hangs, you can force the creation of a dump file. However, you must have previously configured the system to allow the creation of a dump file by:

- Making sure that the system is configured to generate a dump file. Click **Start→Control Panel→System→Advanced→Startup and Recovery→System Failure** to verify the settings.
- Using the registry editor (regedit) to modify the following registry setting:

  Hive: HKEY_LOCAL_MACHINE

  Key: System\CurrentControlSet\Services\i8042prt\Parameters

  Name: CrashOnCtrlScroll

  Data Type:
  REG_DWORD

  Value: 1

  Note: A value of 1 enables the feature.

  To create a dump file, press and hold the right **Ctrl** while pressing **ScrLk** twice from the keyboard. The dump file is generated the next time you power on the machine.

One-button data collection utility

You can use the one–button data collection utility to gather vital product data (VPD) about SAN File System hardware and software. This information can help you analyze a problem as well as collect the data to send to other support personnel.

You can invoke the one-button data collection utility in one of the following ways:

- For an engine in the cluster, perform one of these actions:

  - Start the SAN File System console. Then click **Maintain System→ Collect Diagnostic Data**.
  - Access the engine and from a shell prompt, run **/usr/tank/server/bin/pmf.sh** to collect the default data or add additional parameters to customize the data collection.
- For a client running AIX, access the client and from a shell prompt, run **/usr/tank/client/bin/pmf.sh** to collect the default data or add additional parameters to customize the data collection.
- For a client running Windows, access the client and from a shell prompt, run **C:\Program Files\IBM\Storage Tank\client\bin\pmf.bat** to collect the default data. In addition, you must perform some additional steps to collect all of the data.

Related topics:

- "collectdiag" on page 181
- "pmf" on page 337
- "Software Vital Product Data" on page 47
- "Hardware Vital Product Data"

Hardware Vital Product Data

This section describes the hardware vital product data that is collected by the one–button data collection utility.

Engine:

You can collect hardware information about an engine in the cluster, such as the following:

Table 2. Engine hardware VPD

Hardware area	Component	VPD collected
Processor/PCI devices	Machine	Type, model number, vendor, and serial number.
	Host bridge	Device, vendor, firmware version, and latency.
	ISA bridge	Device, vendor, and firmware version.
	Ethernet controller	For each device: adapter type, vendor, firmware version, latency, and memory usage.
	USB controller	
	VGA controller	
	IDE interface	
	Fibre-channel adapter	
	RSA II adapter	
Memory	Memory	Statistics for total memory available, used, free, shared, buffered, and cached. Additional usage statistics as well.
	Swap space	

Table 2. Engine hardware VPD (continued)

Hardware area	Component	VPD collected
LAN network	Ethernet interfaces	For each device: data received and transmitted, Internet address, network masks, packets, collisions, interrupts, errors, and base memory address.
	Loopback interface	Statistics about Internet address, network masks, packets, collisions, and errors.
	IP routing table	For each destination: gateway address, network masks, flags, and interfaces.
Local storage	SCSI devices	For each device: channel, ID, LUN, vendor, model, version, and type.
	File systems	Device, mount point type, and inodes (total, used, and free).
	Mount points	Device, file system, and read/write settings.

Related topics:

- "One-button data collection utility" on page 45
- "pmf" on page 337

Software Vital Product Data

This section describes the software vital product data that is collected by the one-button data collection utility.

Engine:

You can collect information about the software running on an engine in the cluster, such as the following:

Table 3. Engine software VPD

Software area	Component	VPD collected
Operating system	Machine	Name and version.
	Operating system	Version, build information, and installation date.
	Processes	Owner, ID, binary file, status, runtime parameters, and environment variables. Additional details as well.
	System log files	Collect in their entirety.
	Core files	Operating system core files and corresponding binary file, if present and requested.
Network	Active connections	Protocol (TCP/UDP), local and remote addresses, state, and receive and send queues.
	Active sockets	Type, state, flags, reference count, and full pathname.
	ARP	IP address, hardware address, and device.

Table 3. Engine software VPD (continued)

Software area	Component	VPD collected
Metadata server	Configuration files	Version, Tank.Bootstrap, and Tank.Config files.
	Log files	log.std, log.audit, log.trace, log.cim, log.failover, and log.*.old.
	Core files	Server core file, if present.
	Server configuration	Dump information about: • Version of installed server code • Current server state • Current server role • Protocol (TCP or UDP) • IP address and network mask
	Server state	Dump information about: • Active threads, their state, and activity • Mutexes and current state of each one • Latches and current state of each one • Condition variables and information about each one • Write-ahead log writer thread information.
	Cluster configuration	Dump information about: • Number of servers in cluster • Listing of all servers in cluster • Fileset (container) information • Global disk table and type of each disk (master, system, user) • List of registered clients and information about each one • Heartbeat interval between servers • Current state of High-Availability Manager
Administrative server	Configuration files	tank.properties, cimom.properties, and tank_device_map files.
	Log files	cimom.log, console.log, security.log, and WebSphere Application Server-based trace.log files.
	Core files	Server core file, if present.

Client:

You can collect information about the software running on a client, such as the following:

Table 4. Client software VPD

Software area	Component	VPD collected
Operating system	Machine	Name and version.
	Operating system	Version, build information, and installation date.
	Processes	Owner, ID, binary file, status, runtime parameters, and environment variables. Additional details as well.
	System log files	Collect in their entirety.
	Core files	Operating system core files and corresponding binary file, if present and requested.
SAN File System client	Log files and trace files	All client log and trace files
Network	Active connections	Protocol (TCP/UDP), local and remote addresses, state, and receive and send queues.
	Active sockets	Type, state, flags, reference count, and full pathname.
	ARP	IP address, hardware address, and device.

Related topics:

- "One-button data collection utility" on page 45
- "pmf" on page 337

Chapter 5. Isolating problems with the SAN File System

In most cases, you can use the logs provided by the SAN File System and the symptoms reported by the customer to begin isolating problems.

- For problems that seem to be related to clients, use the logs that are available with the client operating system (such as the system log on AIX clients and the Event Log on Windows clients) to determine the cause of the problem. In addition, you can use the information provided in Chapter 8, "Troubleshooting a SAN File System client", on page 77.
- For problems that seem to be related to administrative user access, use the security log and the administrative log to determine the cause of the problem. If you access these logs through the master Metadata server, you will see a consolidated view of the logs from each of the Metadata servers in the cluster. If you access these logs through a subordinate Metadata server, you will see the logs for that particular Metadata server.

  In addition, you can use the information provided in Chapter 7, "Troubleshooting an Administrative server", on page 67.
- For problems that seem to be related to the cluster, Metadata servers, or metadata, use the server log to determine the cause of the problem.

  In addition, you can use the information provided in Chapter 6, "Troubleshooting the cluster", on page 55. If you access this log through the master Metadata server, you will see a consolidated view of the logs from each of the Metadata servers in the cluster (called the cluster log). If you access these logs through a subordinate Metadata server, you will see the logs for that particular Metadata server.
- For problems that seem to be related to the engines in the SAN File System, you can use the Light Path Diagnostics and the RSA II adapter Web interface to determine the cause of the problem. In addition, you can use the information provided in Chapter 10, "Troubleshooting an engine", on page 87.

In cases where you are not sure whether the problem is related to the SAN File System, SAN, or LAN, use the information in this section to begin isolating the problem.

Note: Certain events, such as operating system reboots or cable disconnects, can cause the SAN File System to lose connectivity to LUNs. If the logs indicates I/O failures for a client or Metadata server, verify the following:

- Configured LUNs are visible from the Metadata servers and SAN File System clients. From the server, you should see both the user LUNs and system LUNs. From a client, you should see only the user LUNs.
- A SAN fabric switch has not lost the zoning configuration. If the operating system on the switch is rebooted, it is possible for the fabric to lose the zoning configuration, which prevents Metadata servers and SAN File System clients from reaching LUNs.

You may need to force the SAN File System to remap LUNs in the event of lost connectivity. Please see your SAN administrator for help with LUN rediscovery options specific to your operating environment.

Identifying SAN problems:

Perform the following steps to determine whether the problem is related to the SAN itself:

1. Determine whether the customer has recently changed the SAN configuration, such as changing the fiber channel cable connections or switch zoning. If so, ask the customer to verify that the changes were correct and if necessary reverse those changes.
2. Verify that all switches and RAID controllers that are used by the SAN File System are powered on and are not reporting any hardware failures. If problems are found, resolve them before proceeding further.
3. Verify that the fiber channel cables that connect the Metadata servers to the switches are securely connected.
4. IBM Subsystem Device Driver (SDD) version 1.4.0 is provided with the SAN File System and provides support for multipath environments. You can use the datapath query commands to view statistics, as well as information about paths and adapters. For information about using the datapath query commands, refer to the *Subsystem Device Driver User's Guide*, which is provided on the SAN File System documentation CD-ROM.
5. If the customer is running a SAN Management tool that you are familiar with and you have access to that tool, use it to view the SAN topology and isolate the failing component. If the customer is not using a SAN Management tool, you can start IBM Tivoli SAN Manager on the master console and use it to view the SAN Topology and isolate the failure. For information about SAN problem determination with IBM Tivoli SAN Manager, contact the Tivoli Storage Area Network (SAN) support center.

Identifying IP networking problems:

Perform the following steps to determine whether the problem is related to the IP network itself:

1. Verify that all switches used by the SAN File System are powered on and are not reporting any hardware failures. If problems are found, resolve them before proceeding further.
2. Verify that the Ethernet cables that connect the Metadata servers to the switches are securely connected.
3. Verify that the Metadata servers, clients, and storage devices are on the same network and subnet.

Identifying storage problems:

Perform the following steps to determine whether the problem is related to the storage devices:

1. Determine whether any other hosts that may be attached to the storage devices are having the same problems.
2. Determine whether a single Metadata server or client is having trouble accessing the storage device or all Metadata servers and clients are experiencing I/O errors.
3. Refer to the documentation for the storage devices for more information about isolating problems with those devices.

Related topics:

- Chapter 6, "Troubleshooting the cluster", on page 55
- "Server diagnostic tools" on page 36

- Chapter 7, “Troubleshooting an Administrative server”, on page 67
- Chapter 8, “Troubleshooting a SAN File System client”, on page 77
- “Client diagnostic tools” on page 42
- Chapter 10, “Troubleshooting an engine”, on page 87
- “Engine diagnostic tools” on page 29

Chapter 6. Troubleshooting the cluster

A SAN File System cluster can contain from two to eight engines, each running a separate instance of a Metadata server. The Metadata servers have one of the following roles:

- **Master**.

  The master Metadata server manages system metadata for the entire cluster. It controls all operations involving system metadata, such as allocation of storage space, coordination of most administrative operations, and access to the global namespace. In addition, the master Metadata server can also perform the same tasks that are performed by subordinate Metadata servers, managing file metadata and workload for one or more filesets.

  One Metadata server at a time can act as the master in a cluster.

- **Subordinate**.

  Subordinate servers manage user metadata and workload for one or more filesets.

  Note: A fileset can be managed by only one Metadata server.

To obtain access to the user data in a specific fileset, clients communicate with the Metadata server that manages that fileset.

Metadata server failures:

When a subordinate Metadata server becomes unresponsive or fails, such as when the operating system crashes or hangs, the engine is automatically restarted. In addition, if you have enabled the automatic restart service (enabled by default), the Metadata server is also automatically restarted.

While the subordinate Metadata server is in the process of restarting, it cannot respond to requests from clients:

- Client requests to the Metadata server and client access of any files served by the Metadata server will fail or be delayed.
- Client applications experience a pause in service while the Metadata server is unavailable (typically, this will last approximately one or two minutes). During this time, active operations of some applications can begin to time out. Whether additional errors occur is based on how the client applications respond to a time-out situation.

When the master Metadata server becomes unresponsive or fails, any clients attempting to access filesets managed by the master experience the same results as clients attempting to access filesets managed by subordinate Metadata servers. In addition, subordinate Metadata servers can be affected by the unavailability of the master. Metadata servers in the cluster rely on a heartbeat mechanism to verify availability. Depending on the length of time that the master Metadata server is unavailable, subordinate Metadata servers may detect the loss of the heartbeat mechanism and cease all activity until the master is available again (or you set a new master).

Related topics:

- Appendix C, “Commands”, on page 157

Troubleshooting a Metadata server

Problem:

A Metadata server has failed. Attempts by the SAN File System to automatically restart the Metadata server have also failed and manual intervention will be required to restart the Metadata server. Clients cannot access file metadata and, in the case of a master Metadata server failure, the cluster itself is no longer available.

Note: If a Metadata server loses power, you will need to reset the RSA II card on the engine hosting the Metadata server to restore proper communication between the RSA II card and the Metadata server. To reset the RSA II card, you will need to access the Web interface for the RSA II card and select the option for resetting the RSA II card.

Investigation steps:

Try the following actions to resolve the problem:

If a subordinate Metadata server has failed, take the following actions:

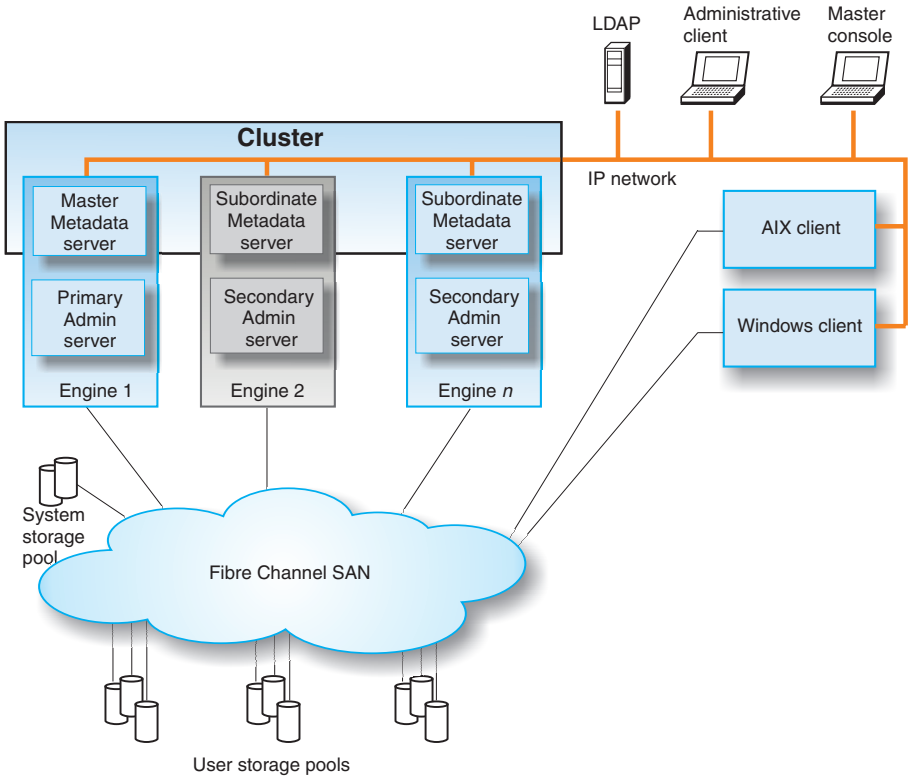

Perform the following steps until the problem is resolved:

1. Use the SAN File System console or the Administrative command-line interface to view the status of the subordinate Metadata server.
2. View the cluster message log to verify that the SAN File System has not been able to restart the subordinate Metadata server. In addition, the messages in this log can provide you with an indication of what the problem may be.

3. Reassign (move) the filesets from the subordinate Metadata server to another Metadata server that is online. The decision to reassign filesets will be based on the length of time it will take to repair the Metadata server.

   Note: Make sure that this Metadata server is actually down and all restart attempts were unsuccessful before reassigning filesets:

   a. Run **lsserver** to verify that the Metadata server is shut down.
   b. Run **lsengine** to verify that the engine hosting the master Metadata server is shut down.
   c. Run **lsserver** to verify that the Metadata server to which you are going to assign the role of master is up and running.

4. Resolve any problems found in the cluster message log that are related to this Metadata server. If the messages indicate a hardware error, use the RSA II Web interface to access the RSA II adapter for the engine hosting the subordinate Metadata server. The RSA II Web Interface can assist you in isolating the hardware problem.
5. If there was an abnormal termination of the Metadata server, you may begin to see errors on the clients, even after the problem with the Metadata server has been resolved by itself with the automatic restart feature or by taking the failed Metadata server offline and reassigning its filesets to another Metadata server. If you begin seeing these types of problems, you will need to restart the affected clients:
   - On clients running the AIX operating system:
     a. Run **rmstclient** to unmount the global namespace, remove the virtual client, and unload the file-system driver.
     b. Run **setupclient** to load the file-system driver, create the virtual client, and mount the global namespace.
   - On clients running the Windows operating system, reboot the system.
6. After repairing the failed Metadata server, bring the server back online (use the **startserver** command from the Administrative command-line interface).
7. If you previously reassigned the filesets for this Metadata server to another server, you can now assign them back to this Metadata server.

If the master Metadata server has failed, take the following actions.

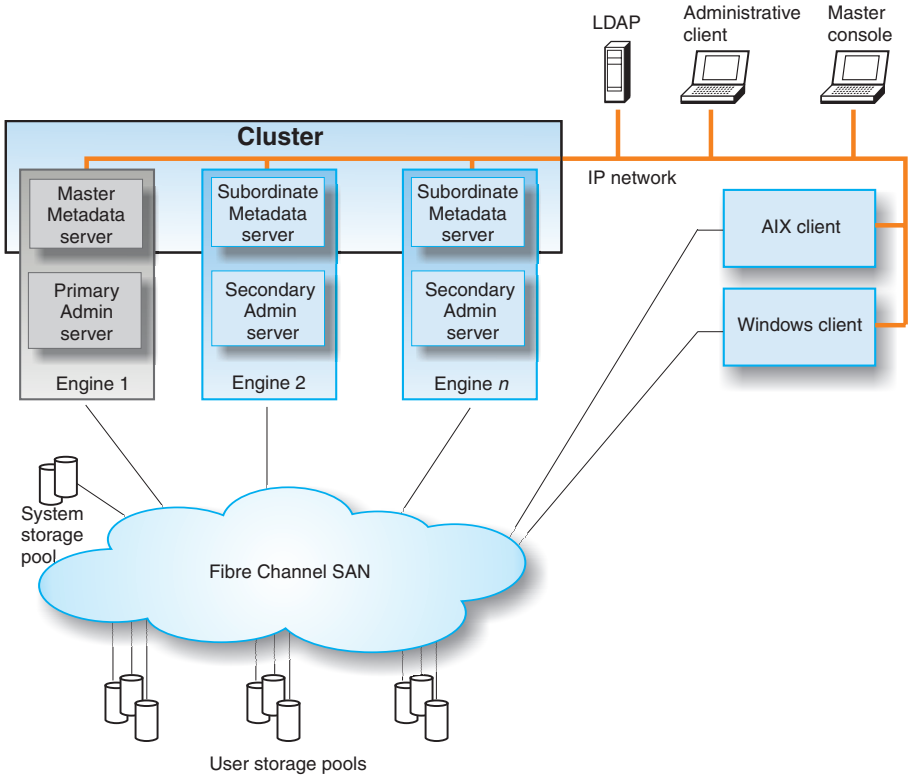

Perform the following steps until the problem is resolved:

1. Use the Administrative command-line interface to view the server message log and verify that the SAN File System has not been able to restart the master Metadata server. In addition, the messages in this log can provide you with an indication of what the problem may be.
2. Define a new master Metadata server for the cluster.

   Note: Make sure that this Metadata server is actually down and all restart attempts were unsuccessful before attempting to set a new master or before reassigning filesets:

   a. Run **lsserver** to verify that the cluster does not have a master Metadata server.

   b. Run **lsengine** to verify that the engine hosting the master Metadata server is shut down.

   c. Run **lsserver** to verify that the Metadata server to which you are going to assign the role of master is up and running.
3. Reassign (move) the filesets from this Metadata server to another Metadata server that is online.
4. Resolve any problems found in the cluster message log that are related to this Metadata server. If the messages indicate a hardware error, use the RSA II Web interface to access the RSA II adapter for the engine hosting the former master Metadata server. The RSA II Web Interface can assist you in isolating the hardware problem.
5. If there was an abnormal termination of the Metadata server, you may begin to see errors on the clients, even after the problem with the Metadata server has been resolved by itself with the automatic restart feature or by taking the failed Metadata server offline and reassigning

its filesets to another Metadata server. If you begin seeing these types of problems, you will need to restart the affected clients:

- On clients running AIX:
  a. Run **rmstclient** to unmount the global namespace, remove the virtual client, and unload the file-system driver.
  b. Run **setupclient** to load the file-system driver, create the virtual client, and mount the global namespace.
- On clients running Windows, reboot the system.

6. After repairing the failed Metadata server, bring the server back online (use the startserver command from the Administrative command-line interface).
7. If you choose, you can now set this Metadata server to be the master Metadata server once again. To set this Metadata server to be the new master, you must first shut down the existing master Metadata server and power off the engine hosting the master.
8. If you previously reassigned the filesets for this Metadata server to another server, you can now assign them back to this Metadata server.

Related topics:

- Chapter 10, "Troubleshooting an engine", on page 87
- "Defining a new master Metadata server" on page 63
- "Reassigning filesets to Metadata servers" on page 64
- "Taking a Metadata server offline" on page 63
- "Administrative commands" on page 158
- "Accessing the RSA II adapter" on page 24

Troubleshooting the local network

Problem:

There is a problem with the local network on which the Metadata servers communicate. The problem may be:

- A network fault. A local network fault can occur if there is a bad Ethernet adapter in an engine or the Ethernet cable is not connected between the Ethernet adapter and the IP network. In the event of a local network fault, the cluster will react as if the Metadata server on which the fault occurred is down.

  The master Metadata server reforms the cluster, excluding the failed Metadata server. The Metadata server itself will go into a wait state, and any filesets assigned to that Metadata server will no longer be available to clients.
- A network partition. A local network partition can occur if there is a problem in the Ethernet network that causes two or more Metadata servers to lose communications with the master Metadata server. The partition containing the master Metadata server will react as if the Metadata servers in the other partition are down. The Metadata servers in the other partition will react as if the master Metadata server is down.

Investigation steps:

Try the following actions to resolve the problem:

If there is a local network fault with one of the subordinate Metadata servers, take the following actions.

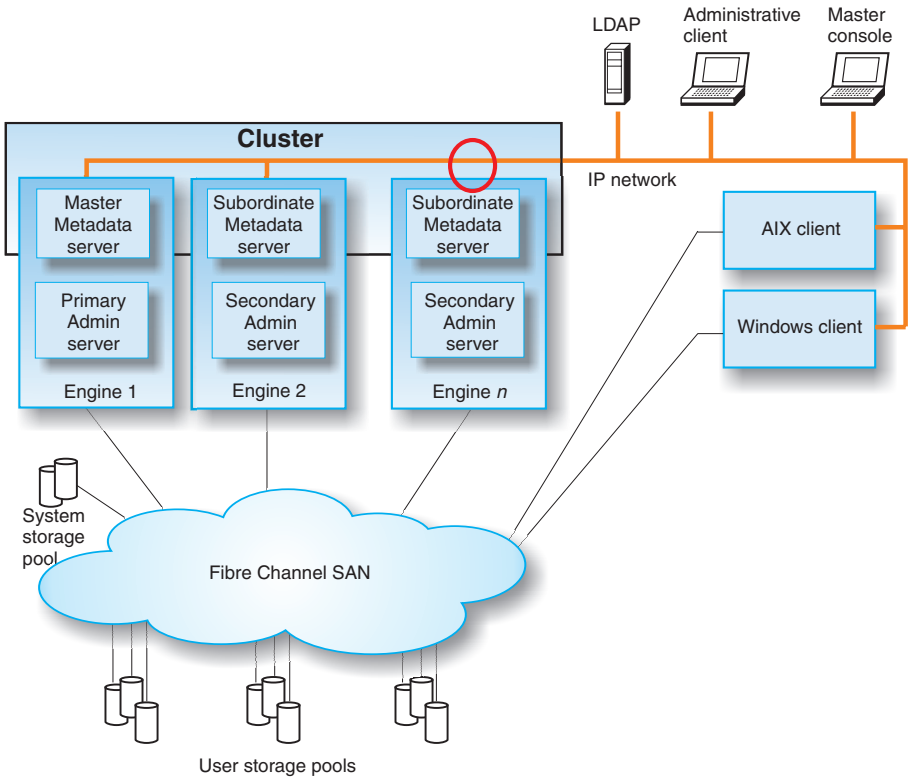

Perform the following steps in order until the problem is resolved:

1. Use the RSA II Web interface to access the RSA II adapter for the engine hosting the subordinate Metadata server.
2. Shut down the engine from the RSA II Web interface. You will not be able to use the Administrative command-line interface or the SAN File System console to shut down the Metadata server because the master Metadata server already considers the server to be shut down.
3. Reassign (move) the filesets from the subordinate Metadata server to another Metadata server that is online. The decision to reassign filesets will be based on the length of time it will take to repair the Metadata server.

   Note: Make sure that this Metadata server is actually down and all restart attempts were unsuccessful before reassigning filesets.

   a. Run **lsserver** to verify that the Metadata server is shut down.
   b. Run **lsengine** to verify that the engine hosting the Metadata server is shut down.
   c. Run **lsserver** to verify that the Metadata server to which you are going to assign the filesets is up and running.
4. After repairing the network fault, you can have the Metadata server rejoin the cluster:
   a. Start the Metadata server.
   b. Wait for the cluster to be reformed to include this Metadata server.
   c. Use the SAN File System console or the Administrative command-line interface to verify that all Metadata servers in the cluster are in an Online state.

5. If you previously reassigned the filesets for this Metadata server to another server, you can now assign them back to this Metadata server.
   a. Run **lsserver** to verify that the Metadata server from which you will move the filesets is shut down.
   b. Run **lsengine** to verify that the engines hosting the Metadata servers are shut down.
   c. Run **lsserver** to verify that the Metadata server to which you are going to assign the filesets is up and running.

If there is a network partition, take the following actions.

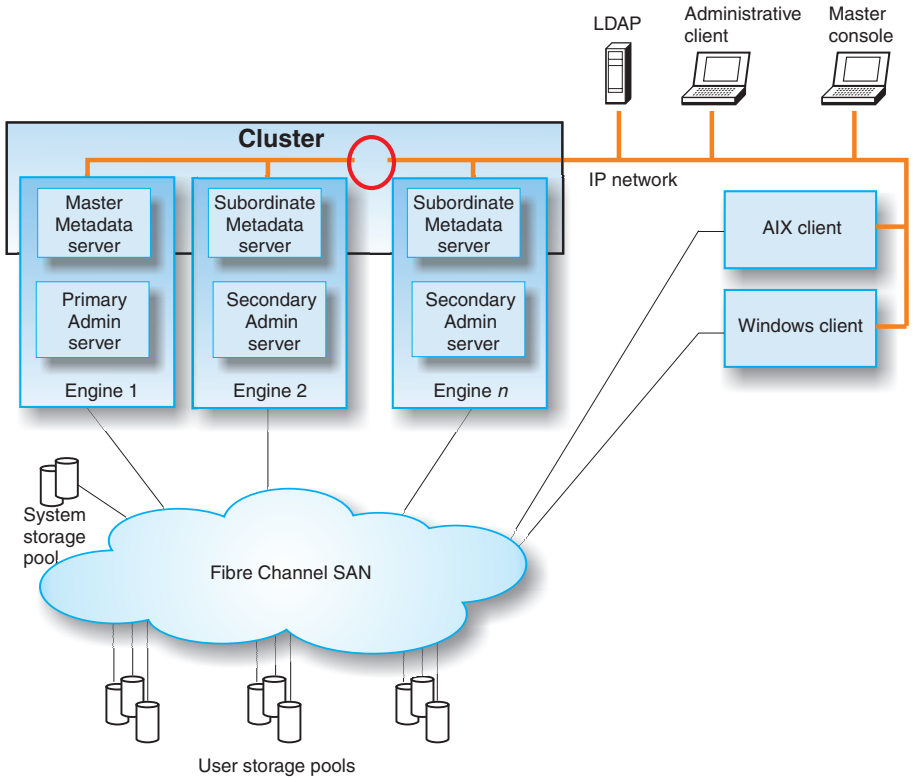

Perform the following steps in order until the problem is resolved:

1. Use the RSA II Web interface to access the RSA II adapter for each of the partitioned engines.
2. Shut down each engine from the RSA II Web interface. You will not be able to use the Administrative command-line interface or the SAN File System console to shut down the Metadata server because the master Metadata server already considers the servers to be shut down.
3. Reassign (move) the filesets from the subordinate Metadata server to another Metadata server that is online. The decision to reassign filesets will be based on the length of time it will take to repair the Metadata server.

   Note: Make sure that this Metadata server is actually down and all restart attempts were unsuccessful before reassigning filesets.
   a. Run **lsserver** to verify that the Metadata server is shut down.
   b. Run **lsengine** to verify that the engine hosting the Metadata server is shut down.
   c. Run **lsserver** to verify that the Metadata server to which you are going to assign the filesets is up and running.

4. After repairing the network partition, you can have the Metadata servers rejoin the cluster:
   a. Start each Metadata server.
   b. Wait for the cluster to be reformed to include these Metadata servers.
   c. Use the SAN File System console or the Administrative command-line interface to verify that all Metadata servers in the cluster are in an Online state.
5. If you previously reassigned the containers for these Metadata servers to other servers, you can now assign them back to these Metadata servers.
   a. Run **lsserver** to verify that the Metadata server from which you will move the filesets is shut down.
   b. Run **lsengine** to verify that the engines hosting the Metadata servers are shut down.
   c. Run **lsserver** to verify that the Metadata server to which you are going to assign the filesets is up and running.

Related topics:

- Chapter 10, "Troubleshooting an engine", on page 87
- "Defining a new master Metadata server" on page 63
- "Reassigning filesets to Metadata servers" on page 64
- "Taking a Metadata server offline" on page 63
- "Administrative commands" on page 158

Resolution Procedures

You can use the procedures in this section to help you resolve problems with the Metadata server. These procedures include:

- "Shutting down an engine from the RSA Web interface"
- "Taking a Metadata server offline" on page 63
- "Defining a new master Metadata server" on page 63
- "Reassigning filesets to Metadata servers" on page 64
- "Bringing a Metadata server online" on page 64

Shutting down an engine from the RSA Web interface

Steps:

Perform these steps to shut down an engine from the RSA Web interface. These steps assume that you are attempting to access the Web interface from the master console and that you have already obtained an RSA user account.

1. Open a browser and point it to the URL for the RSA adapter that is located in the engine that you want to shut down.
2. From the Enter Network Password panel, type your RSA user name and password. Then click **OK**.
3. From the Welcome panel, select a session timeout value and click **Continue**.
4. From the left navigation pane of the System Health panel, click **Server→Tasks→Power/Restart**.
5. From the Power/Restart panel, Click **Power off server immediately**.

6. From the left navigation pane of the System Health panel, click **Sign off**.

Related topics:

- "Accessing the RSA II adapter" on page 24
- "Administrative commands" on page 158

Taking a Metadata server offline

Steps:

Perform these steps to take a subordinate Metadata server offline. These steps assume that you are using the Administrative command-line interface.

1. Run the **stopserver** command to stop the Metadata server.
2. Reassign filesets that are currently assigned to the Metadata server.

Perform these steps to remove a master Metadata server from the cluster. These steps assume that you are using the Administrative command-line interface.

1. Run the **stopserver** command to stop the master Metadata server.
2. Run the **setmaster** command from a subordinate Metadata server to set a new master.
3. Reassign filesets that were assigned to the Metadata server that you just stopped.

Result:

The Metadata server is now offline.

Related topics:

- "Reassigning filesets to Metadata servers" on page 64
- "Bringing a Metadata server online" on page 64

Defining a new master Metadata server

Steps:

Perform these steps to define a new master Metadata server. These steps assume that you are using the Administrative command-line interface.

1. Verify that the current master Metadata server is offline.
   a. Run **lsserver** to verify that the cluster does not have a master Metadata server.
   b. Run **lsengine** to verify that the engine hosting the master Metadata server is shut down.
   c. Run **lsserver** to verify that the Metadata server to which you are going to assign the role of master is up and running.
2. Access the subordinate Metadata server that you want to set as the new master Metadata server.
3. Run the **setmaster** command to set it as the new master Metadata server.
4. If the old master Metadata server was managing any filesets, reassign those filesets to other Metadata servers in the cluster.

Related topics:

- "Reassigning filesets to Metadata servers"

Reassigning filesets to Metadata servers

Steps:

Perform these steps to assign a fileset to a new Metadata server. These steps assume that you are using the Administrative command-line interface.

1. Run **lscontainer** to list all of the filesets assigned to a specific Metadata server.
   ```
   tanktool lscontainer -server metadata_server_name
   ```
2. Run **setcontainerserver**, specifying the new Metadata server to which the filesets are to be assigned.

   Note: You can assign multiple filesets to a Metadata server from a single **setcontainerserver** command. However, you will need to run the **setcontainerserver** command for each Metadata server to which you intend to assign filesets.

Related topics:

- "Administrative commands" on page 158

Bringing a Metadata server online

Steps:

Perform these steps to bring a Metadata server online. These steps assume that you are using the Administrative command-line interface.

1. Run the **startserver** command to start the Metadata server.
2. Run the **stopserver** command to stop the Metadata server to which the filesets are currently assigned.
3. Reassign any filesets that need to be assigned to this Metadata server.
4. Run the **startserver** command to start the Metadata server that you previously stopped.

Result:

The Metadata server is now online and managing filesets.

Related topics:

- "Reassigning filesets to Metadata servers"
- "Taking a Metadata server offline" on page 63

Repairing metadata

Context:

The metadata checker utility should be run in the following situations:

- Periodically, to validate the integrity of SAN File System metadata
- After reverting a fileset to a FlashCopy image
- Following the failure of a Metadata server to ensure that no metadata corruption has occurred and to have the metadata checker utility attempt to repair any inconsistencies that are found

- In response to error messages in the cluster log that indicate possible metadata inconsistency
- In response to system behavior that implies possible metadata inconsistency

Steps:

Perform these steps to run the metadata checker utility:

1. Use either the Administrative command-line interface or the SAN File System console to access the master Metadata server.
   - If you use the Administrative command-line interface, run the **startmetadatacheck** command with the check option.
   - If you use the SAN File System console, click **Maintain system→Check metadata** to display the Check Metadata panel. Choose the type of metadata to check and the location of the metadata to be checked.
2. After verifying the extent of the corruption, you can run the metadata checker utility again with the repair option specified.

Related topics:

- "Administrative commands" on page 158
- "reverttoimage" on page 249
- "startmetadatacheck" on page 274
- "stopmetadatacheck" on page 292

Backing up system metadata

Steps:

Perform these steps to back up system metadata:

1. From the Administrative command-line interface, create a new system-metadata disaster-recovery dump file.

   ```
   tanktool> mkdrfile dr_file_1
   ```

   A message confirming the creation of the file is displayed.
2. Verify that the system-metadata disaster-recovery file was created.

   ```
   tanktool> lsdrfile
   ```

   A list of all of the system-metadata disaster-recovery files is displayed, including the file just created (dr_file_1).
3. Exit the Administrative command-line interface and change directories to the directory where the recovery files are stored. List the files in this directory to verify that you have at least four files.

   ```
   tanktool>exit
   # cd /usr/tank/server/DR
   # ls
   ```

   You should see at least four files. One will be the recovery file that you created (with a dump extension). You will also see three script files (all three files begin with TankSysCLI).
4. Save these files with your normal file data backup procedures. These files are used to recover system metadata as part of the disaster recovery procedures.

Related topics:

- "Administrative commands" on page 158
- "Restoring SAN File System metadata" on page 149

Chapter 7. Troubleshooting an Administrative server

Each Metadata server in the SAN File System cluster runs an instance of the Administrative server, which provides administrative access to the Metadata server. The Administrative server running on the engine that hosts the master Metadata server is referred to as the primary Administrative server. All other Administrative server instances are referred to as secondary Administrative servers.

Administrative server components:

The Administrative server contains three main components:

- SAN File System console – a set of servlets that run on WebSphere® Application Server. The console provides a Web browser interface to the SAN File System. Users access the console through a secure connection by going to:

  https://*master_metadata_server_IP_address*:7979/tank

  If users point to the IP address of a subordinate Metadata server and the master Metadata server is online, they will automatically be redirected to the IP address of the master Metadata server.
- Administrative command-line interface (CLI) – the program (called tanktool) that is available on each engine in the cluster. To access the Administrative CLI, users must initiate a secure shell (SSH) session with the master Metadata server.

  Users who initiate an SSH session with a subordinate Metadata server are not automatically redirected to the master. However, most Administrative CLI commands must be run from the master Metadata server, so users should typically initiate the SSH session with the master.

  Note: One exception to running a command from the master Metadata server is the **setmaster** command, which is used to designate a new master Metadata server in the event that there is a failure of the current master. This command is run from the subordinate Metadata server that is going to become the master.

  Many commands will provide output if run from a subordinate Metadata server. However, the output may not be what you expect. For example, the **lsserver** command provides information about the Metadata servers in the cluster if you run the command from the master Metadata server. If you run this command from a subordinate Metadata server, you will see details about that specific Metadata server only.
- Administrative agent – part of the SAN File System implementation of the Common Information Model (CIM), which is a model for describing management and information interchange between agents and managers. The Administrative agent is used for all native SAN File System operations, such as:
  - Interfacing with the Lightweight Directory Access Protocol (LDAP) server for user authentication
  - Automatically attempting to restart the Metadata server in the event of a failure
  - Registering with the service location protocol (SLP), which is a CIM agent directory service

Startup sequence:

If you have enabled the automatic restart service, the Administrative agent performs these activities when it starts up:

1. Connects to the local Metadata server.
2. Learns the location of the master Metadata server.
3. Connects to the LDAP server.
4. Registers with the SLP daemon.

Related topics:

• "Server diagnostic tools" on page 36

Troubleshooting user access to the console

Problem:

A user was unable to access the SAN File System console from a Web browser.

Investigation steps:

Try the following actions to resolve the problem:

If the user received access denied or unauthorized user errors.

Perform the following steps in order until the problem is resolved:

1. Attempt to sign on using a user name and password that you know are valid.
   a. If you can sign on using a different user name and password, verify that the "unauthorized" user's user name and password are valid. The user name and password must be created on the Lightweight Directory Access Protocol (LDAP) server.
      - You can run the **lsadmuser** command to list all of the administrative users. If the user is not listed, you can run the **ldapsearch** command from the bash shell to determine if the user is defined in the LDAP server (see the help that is available with **ldapsearch** to learn more about that command). In addition, you can run **ldapsearch** - help to view information about the **ldapsearch** command online.
      - Use the documentation that is provided with the LDAP server to verify that the account for the user was set up correctly with the LDAP server.

   b. Attempt to use this user's user name and password with the other administrative access method. For example, if the user received the error while using the SAN File System console, attempt to sign on to the Administrative command-line interface. If the user name and password work with the other administrative access method, suspect a problem with the previous administrative access method (the console or the command-line interface).
2. Verify that the LDAP server is running and that there is no configuration problem between the LDAP server and the Administrative agent. Check the Administrative log to determine whether you are receiving errors about the SAN File System not being able to connect to the LDAP server.

If the user received page not found errors.

Perform the following steps until the problem is resolved:

1. Verify that you entered the correct URL for the primary Administrative server.
2. Verify that you can access the engine hosting the primary Administrative server.
   a. From a shell prompt, attempt to ping the engine. If you cannot ping the engine, suspect either:
      - IP network problem
      - Hardware problem on the engine (see Chapter 10, "Troubleshooting an engine", on page 87)
3. Verify that the Administrative agent is running.

If the user received Administrative agent or CIM agent not found errors.

Perform the following steps in order until the problem is resolved:

1. Run the **ps** shell command to verify that the Administrative agent is running.

   ```
   ps -ef w | grep -i cimom.cimom
   ```

   You should see results similar to this:

   ```
   root      3822     1  0 07:26 ?        S      0:00
   /bin/bash /etc/rc.d/init.d/cimom
   start /usr/tank/admin
   root      4070  3822  0 07:26 ?        S      0:08
   /opt/IBMJava2-131/jre/bin/exe/java -Xms128m -Xmx256m
   com.ibm.cimom.CIMOM -RootDir
   root      4087  4070  0 07:26 ?        S      0:00
   /opt/IBMJava2-131/jre/bin/exe/java -Xms128m -Xmx256m
   com.ibm.cimom.CIMOM -RootDir
   root      4088  4087  0 07:26 ?        S      0:00
   /opt/IBMJava2-131/jre/bin/exe/java -Xms128m -Xmx256m
   com.ibm.cimom.CIMOM -RootDir
   root      4089  4087  0 07:26 ?        S      0:00
   /opt/IBMJava2-131/jre/bin/exe/java -Xms128m -Xmx256m
   com.ibm.cimom.CIMOM -RootDir
   root      4090  4087  0 07:26 ?        S      0:00
   /opt/IBMJava2-131/jre/bin/exe/java -Xms128m -Xmx256m
   com.ibm.cimom.CIMOM -RootDir
   root      4098  4087  0 07:27 ?        S      0:00
   /opt/IBMJava2-131/jre/bin/exe/java -Xms128m -Xmx256m
   com.ibm.cimom.CIMOM -RootDir
   ldap      4099   758  0 07:27 ?        S      0:00
   /usr/sbin/slapd -u ldap
   root      4100   946  0 07:27 ?        S      0:00
   /opt/was/java/jre/bin/exe/java
   -Xbootclasspath/p:/opt/was/java/jre/lib/ext/ibmorb
   root      4101  4087  0 07:27 ?        S      2:11
   /opt/IBMJava2-131/jre/bin/exe/java -Xms128m -Xmx256m
   com.ibm.cimom.CIMOM -RootDir
   root      4102  4087  0 07:27 ?        S      0:00
   /opt/IBMJava2-131/jre/bin/exe/java -Xms128m -Xmx256m
   com.ibm.cimom.CIMOM -RootDir
   ```

2. If the Administrative agent is not running, run /usr/tank/admin/bin/startCimom to start the Administrative agent. If you run into problems starting the Administrative agent, view the Administrative log (/usr/tank/admin/cimom.log) or /usr/tank/admin/log/stderr.log and attempt to resolve any problems that you find.
3. Verify that the console has been started.

   ```
   /opt/was/bin/serverStatus.sh metadata_server
   ```

where *metadata_server* is the host name of the server that you are trying to access.

Note: Run the script with no parameters to obtain additional help on using the script.

4. If the console is not running, run /usr/tank/admin/bin/startConsole to start the console. If you run into problems starting the console, view the log indicated by the error message and attempt to resolve any problems that you find.
5. If both the Administrative agent and console are running, suspect an IP networking problem.

Related topics:

- "Troubleshooting a Metadata server" on page 56
- Chapter 10, "Troubleshooting an engine", on page 87
- "startCimom" on page 343

Troubleshooting user access to the Administrative command-line interface

Problem:

A user was unable to access the Administrative command-line interface.

Investigation steps:

Try the following actions to resolve the problem:

If the user received access denied errors

Perform the following steps until the problem is resolved:

1. Using SSH, access the engine that is having failures.
2. View the tank.passwd file to verify that the user name and password are correct.

   ```
   # cat ~/.tank.passwd
   ```

3. Re-create the password file.

   ```
   cd ~; /usr/tank/admin/bin/tankpasswd -u ldap_username -p ldap_passwd
   ```

   where the *ldap_username* and *ldap_password* are for the user that is having problems.
4. Run the **lsserver** command to verify that you have access to the Metadata server.

   ```
   tanktool>lsserver
   ```

5. Verify that the LDAP server is running and that there is no configuration problem between the LDAP server and the Administrative agent. Check the Administrative log to determine whether you are receiving errors about the SAN File System not being able to connect to the LDAP server.

If the user received Administrative agent or CIM agent not found errors.

Perform the following steps until the problem is resolved:

1. Run the **ps** shell command to verify that the Administrative agent is running.

```
ps -efw | grep -i cimom.cimom
```

You should see results similar to this:

```
root      4070  3822  0 07:26 ?        S      0:08
/opt/IBMJava2-131/jre/bin/exe/java -Xms128m -Xmx256m
com.ibm.cimom.CIMOM -RootDir
root      4087  4070  0 07:26 ?        S      0:00
/opt/IBMJava2-131/jre/bin/exe/java -Xms128m -Xmx256m
com.ibm.cimom.CIMOM -RootDir
root      4088  4087  0 07:26 ?        S      0:00
/opt/IBMJava2-131/jre/bin/exe/java -Xms128m -Xmx256m
com.ibm.cimom.CIMOM -RootDir
root      4089  4087  0 07:26 ?        S      0:00
/opt/IBMJava2-131/jre/bin/exe/java -Xms128m -Xmx256m
com.ibm.cimom.CIMOM -RootDir
root      4090  4087  0 07:26 ?        S      0:00
/opt/IBMJava2-131/jre/bin/exe/java -Xms128m -Xmx256m
com.ibm.cimom.CIMOM -RootDir
root      4098  4087  0 07:27 ?        S      0:00
/opt/IBMJava2-131/jre/bin/exe/java -Xms128m -Xmx256m
com.ibm.cimom.CIMOM -RootDir
ldap      4099   758  0 07:27 ?        S      0:00
/usr/sbin/slapd -u ldap
root      4101  4087  0 07:27 ?        S      2:11
/opt/IBMJava2-131/jre/bin/exe/java -Xms128m -Xmx256m
com.ibm.cimom.CIMOM -RootDir
root      4102  4087  0 07:27 ?        S      0:00
/opt/IBMJava2-131/jre/bin/exe/java -Xms128m -Xmx256m
com.ibm.cimom.CIMOM -RootDir
```

2. If the Administrative agent is not running, use the **startCimom** command to start the Administrative agent. If you run into problems starting the Administrative agent, view the Administrative log or /usr/tank/admin/log/stderr.log and attempt to resolve any problems that you find.
3. Verify that the console has been started.

   ```
   /opt/was/bin/serverStatus.sh metadata_server
   ```

   where *metadata_server* is the host name of the server that you are trying to access.

   Note: Run the script with no parameters to obtain additional help on using the script.
4. If the console is not running, run **/usr/tank/admin/bin/startConsole** to start the console. If you run into problems starting the console, view the log indicated by the error message and attempt to resolve any problems that you find.
5. If both the Administrative agent and console are running, suspect an IP networking problem.

Related topics:

- "Accessing an engine through SSH" on page 23
- "lsserver" on page 221

Troubleshooting user access to a specific task or command

Problem:

A user does not see a specific task listed from the SAN File System console or a user cannot run a specific command from the Administrative command-line interface. The user receives authorization errors.

Investigation steps:

Try the following actions to resolve the problem:

Within the SAN File System, authorization to perform tasks or run commands is based on the user role, which is defined in the LDAP server database. There are four valid roles in the SAN File System, ranging from Monitor, which provides a basic level of access, up to Administrator, which provides full access to the system. These roles determine what commands and tasks can be performed.

To resolve this problem, verify that the user's role is sufficient to perform the specified task or run the specified command. You can use the **lsadmuser** command to verify the roles for each administrative user. If the role is not sufficient, either update the user's role to a higher level of access or use a different user name and password (one that has sufficient authorization) to access the SAN File System.

Note: Use the documentation that came with the LDAP server to understand how to modify a role for a user.

Related topics:

- Appendix C, "Commands", on page 157

Resolution procedures

You can use the procedures in this section to help you resolve problems with the Administrative server. These procedures include:

- "Verifying that the Administrative agent is running"
- "Verifying that the console is running" on page 73
- "Replacing expired LDAP and CIMOM certificates" on page 73
- "Configuring LDAP for SAN File System" on page 74

Verifying that the Administrative agent is running

Steps:

Perform these steps to verify that the CIM agent is running:

1. Access the engine hosting the master Metadata server (by using SSH or the RSA II remote console interface).
2. From a bash shell prompt, enter the following command:

```
ps -efw | grep -i cimom.cimom
```

Result:

If the Administrative agent is running, you will see output similar to the following:

```
root      4070  3822  0 07:26 ?        S      0:08
/opt/IBMJava2-131/jre/bin/exe/java -Xms128m -Xmx256m
com.ibm.cimom.CIMOM -RootDir
root      4087  4070  0 07:26 ?        S      0:00
/opt/IBMJava2-131/jre/bin/exe/java -Xms128m -Xmx256m
com.ibm.cimom.CIMOM -RootDir
```

```
root      4088  4087  0 07:26 ?        S      0:00
/opt/IBMJava2-131/jre/bin/exe/java -Xms128m -Xmx256m
com.ibm.cimom.CIMOM -RootDir
root      4089  4087  0 07:26 ?        S      0:00
/opt/IBMJava2-131/jre/bin/exe/java -Xms128m -Xmx256m
com.ibm.cimom.CIMOM -RootDir
root      4090  4087  0 07:26 ?        S      0:00
/opt/IBMJava2-131/jre/bin/exe/java -Xms128m -Xmx256m
com.ibm.cimom.CIMOM -RootDir
root      4098  4087  0 07:27 ?        S      0:00
/opt/IBMJava2-131/jre/bin/exe/java -Xms128m -Xmx256m
com.ibm.cimom.CIMOM -RootDir
ldap      4099   758  0 07:27 ?        S      0:00
/usr/sbin/slapd -u ldap
root      4101  4087  0 07:27 ?        S      2:11
/opt/IBMJava2-131/jre/bin/exe/java -Xms128m -Xmx256m
com.ibm.cimom.CIMOM -RootDir
root      4102  4087  0 07:27 ?        S      0:00
/opt/IBMJava2-131/jre/bin/exe/java -Xms128m -Xmx256m
com.ibm.cimom.CIMOM -RootDir
```

If the Administrative agent is not running, you will not see any output.

Related topics:

- "Accessing an engine through SSH" on page 23
- "startCimom" on page 343

Verifying that the console is running

Steps:

Perform these steps to verify that the console is running:

1. Access the engine hosting the master Metadata server by using SSH.
2. From a shell prompt, enter the following command:
   `/opt/was/bin/serverStatus.sh` *`metadata_server`*

   where *metadata_server* is the server you are trying to access.

Result:

If the console is not running, you will see error messages.

Related topics:

- "Accessing an engine through SSH" on page 23
- "startCimom" on page 343

Replacing expired LDAP and CIMOM certificates

Context:

CIMOM and LDAP certificates can expire. When this happens, they must be replaced. If you get an error saying: "Invalid key in truststore," you must update your LDAP certificate.

Steps:

Perform the following steps to replace an expired certificate.

1. Obtain the current certificate. LDAP certificates are obtained from the LDAP administrator. CIMOM certificates are created by the **mktruststore** command. See step 4.
2. On each engine, run **stopConsole**, then **stopCimom**.
3. On the master engine, change to /usr/tank/admin.
4. Run **bin/mktruststore**. As a parameter, use the path and file name of the LDAP certificate, if it exists.
5. Use **scp** to copy the truststore to each engine in the cluster.

   Note: Do not run the **mktruststore** command on each engine. You must copy the truststore to each engine.
6. On each engine, run **/usr/tank/admin/bin/startCimom**. Then run **/usr/tank/admin/bin/startConsole**.
7. If needed, you can now extract the CIMOM certificate for your third-party CIM application.

Related topics:

- "mktruststore" on page 335

Configuring LDAP for SAN File System

There are several LDAP configuration settings in the cimom.properties file that must be set up during configuration.

Steps:

Perform the following steps to configure LDAP:

1. Change to the /usr/tank/admin/config/ directory.
2. Open the cimom.properties file.
3. Modify the following parameters:

Table 5. LDAP parameters

Parameter name:	Sample values:	Description:
Port	example: 5989	The port on which CIMOM will listen. Set this to 5989.
TrustPassword		The password used when configuring the truststore.
Authorization		Set to false if you want everyone to be able to access the CLI without access controls. If this is set to false, the GUI is unusable.
ldap.cache.age=600		Maximum age of items in the LDAP Cache. Use 0 to disable the cache.
userrole.ldap.location		The IP address of the LDAP server.
userrole.ldap.bind.dn	example: cn=root	The distinguished name of an authorized LDAP user.
userrole.ldap.cred	example: password	The password of the LDAP user.

Table 5. LDAP parameters (continued)

Parameter name:	Sample values:	Description:
userrole.ldap. secured.connection		The flag to enable secured LDAP communication. Set to true, indicates that it uses SSL; set to false, indicates that it uses an open socket.
userrole.ldap.version		Set to 2 if the LDAP server does not support v3.
userrole.ldap.insecured.port	example: 389	The port on which the LDAP server should be listening for an insecure connection. 389 is standard.
userrole.ldap.secure.port	example: 636	The port on which the LDAP server should be listening for a secure connection. 636 is standard.
ldap.basedn.roles	example: ou=Roles,o=ibm,c=us	The base distinguished name to search for roles. This is the location in the LDAP hierarchy to find the role definitions.
ldap.basedn.users	example: ou=Users,o=ibm,c=us	The base distinguished name to search for users. This is the location in the LDAP hierarchy to find users.
ldap.user.filter	example: (&(uid=%v) (objectClass=inetOrgPerson))	The search filter to find a user.
ldap.user.id.attr	example: uid	The attribute that holds the User ID in the user's objectClass.
ldap.role.filter	example: (&(cn=%v) (objectClass=accessRole))	The role filter to find a role.
ldap.role.id.attr	example: cn	The attribute that holds the name of a role in the role's objectClass.
ldap.role.mem.id.attr	example: member	The attribute that holds the members of a role in the role's objectClass.
LogOnly		Setting this to true ensures that stdout.log in /usr/tank/admin/log is not a copy of cimom.log in the same place. This is recommended.
Note: Default values for parameters not listed in this table are acceptable.		

Chapter 8. Troubleshooting a SAN File System client

SAN File System clients run the following operating systems:

- AIX 5.1 (32–bit only) – maintenance level 3 or higher
- Windows 2000 Server or Advanced Server – Service Pack 3

Both clients see the same global namespace. File permissions are based on the operating system on which the files were created:

- When files are created and accessed by Windows clients, all Windows security features are available and enforced.
- When files are created and accessed by AIX clients, all AIX security features are available and enforced.
- When files created by an AIX client are accessed by a Windows client, access is controlled by the permissions of the Other permission bits.
- When files created by a Windows client are accessed by an AIX client, access is controlled by the permissions of the Everyone group.

Use the following topics to determine problem areas within SAN File System clients:

- "Troubleshooting client access to data"
- "Troubleshooting client performance problems" on page 80

Troubleshooting client access to data

Problem:

A client cannot access or update user data.

Investigation steps:

Try the following actions to resolve the problem:

If a client cannot create a new file

Perform the following steps until the problem is resolved:

1. Verify that the master Metadata server is online.
2. Verify that the client has connectivity to the master Metadata server.
   - From the client, attempt to ping or establish an SSH session with the Metadata server.
   - From the Administrative command-line interface, run the **lsclient** command to see a list of all clients currently being served by the Metadata servers in the cluster.

     tanktool> lsclient
3. Use the **ls -l** command to list the directory in which the file is to be created to verify that the entire path is correct and accessible.
4. Verify that you are logged into the client with a user name that has authorization to write files to directories.
   - From an AIX client, use the **id** command.

- From a Windows client, right-click the directory. Then click **Properties→Security**.

5. Verify that your user name has the authorization to write files in the specific directory.

   Note: If you are logged into an AIX client as root or a Windows client as Administrator, make sure that you are running on a privileged client.

6. Check the quota of the fileset in which the parent directory resides to ensure that there is sufficient space to accommodate the new file and that the fileset is attached.
   a. Access the master Metadata server through either the SAN File System console or the Administrative command-line interface.
   b. List the containers to view the server to which the fileset is attached, the quota percentage and type, the attach point, and the directory.
      - From the Administrative command-line interface, run the **tanktool lscontainer -l** command.
      - From the SAN File System, click **Manage Filing→Containers**.
7. Make sure that the storage pool in which the file will be stored has sufficient space to store the file.
8. Verify that the client can access the storage device.
   - Use the **datapath query** command to ensure that the operating system can access the storage device.
   - On an AIX client, you can use the **stfsdisk** command to determine which disks can be seen. Then you can use the **lsvol -l** command to correlate the disk back to the device.

If a client cannot access an existing file

Perform the following steps until the problem is resolved:

1. Verify that your user name has the authorization to read files and that it has authorization to read the specific file.
2. Verify that the client can access the storage device.
   - Use the **datapath query** command to ensure that the operating system can access the storage device.
   - On an AIX client, you can use the **stfsdisk** command to determine which disks can be seen. Then you can use the **lsvol -l** command to correlate the disk back to the device.
3. Suspect a problem with corrupt SAN File System metadata.

If a client cannot update an existing file

Perform the following steps until the problem is resolved:

1. Verify that your user name has the authorization to write to the specific file.
2. Verify that the client can access the storage device.
   - Use the **datapath query** command to ensure that the operating system can access the storage device.
   - On an AIX client, you can use the **stfsdisk** command to determine which disks can be seen. Then you can use the **lsvol -l** command to correlate the disk back to the device.
3. Suspect a problem with corrupt SAN File System metadata.

If a client cannot access any data

Perform the following steps until the problem is resolved:

1. Verify that the client can access the cluster.

   From the client, attempt to sign on to the SAN File System console. If you cannot, suspect one of the following:

   - One or more Metadata servers in the cluster are down. See Chapter 6, "Troubleshooting the cluster", on page 55.

     Note: The Administrative agent will automatically attempt to restart a Metadata server if it goes down for any reason. Therefore, if you cannot access a Metadata server, you might want to wait a few minutes and try again just to ensure that it is not in the process of restarting.

   - There is an IP network problem between the client and the cluster. See Chapter 5, "Isolating problems with the SAN File System", on page 51.

2. View the system logs and look for any errors that may indicate I/O errors. Attempt to resolve all of these errors.

   - From a client running Windows, you can use the Event Viewer to view the Event Log.
   - From a client running AIX, you can view logging and tracing output, assuming it was previously enabled with the syslog facility.

3. From the Administrative command-line interface on the master Metadata server, run the **lslun** command to verify that the luns are available.

   If the disks are not available, you can rediscover all disks by running the following commands from a Metadata server:

   a. Run **stopserver** from the Administrative command-line interface to stop the SAN File System.
   b. Run **rmmod qla2300**.
   c. Run **insmod qla2300** .
   d. Run **/etc/rc.d/init.d/sanfs start** to start the SAN File System.
   e. Run **startserver** from the Administrative command-line interface to start the Metadata server.

4. Verify that the client can access the storage device:

   - Use the **datapath query device** command to ensure that the operating system can access the storage device.
   - On a client running AIX, use the **stfsdisk** command to determine which disks can be seen. Then, from the Administrative command-line interface, you can use the **lsvol -l** command to correlate the disk back to the device.

   To rediscover all disks from a client running AIX, run the client command **stfsdisk -discover**.

   The disks should automatically be rediscovered on a client running Windows. However, it they are not:

   a. Right-click My Computer.
   b. Select Manage.
   c. Select Storage\Disk Management.
   d. From the Action menu, select Rescan Disks.

5. Verify that none of the Metadata servers in the cluster has recently terminated abnormally. If there was an abnormal termination of a Metadata server, you may begin to see errors on the clients even after problems with the failed Metadata server have been resolved. If you begin seeing these types of problems, you will need to restart clients:
   - On clients running AIX:
     a. Run **rmstclient** to unmount the global namespace, remove the virtual client, and unload the file-system driver.
     b. Run **setupclient** to load the file-system driver, create the virtual client, and mount the global namespace.
   - On clients running Windows, reboot the system.
6. If the client cannot access user data, suspect the SAN route from the client. See Chapter 5, "Isolating problems with the SAN File System", on page 51.
7. Suspect a problem with corrupt SAN File System metadata.

If a client running Windows receives delayed write failure errors

A delayed write failure error may appear as a message box on the client desktop or in the Event Log. This message indicates that there has been an error writing data from the local file system cache to a storage device. Perform these steps until the problem is resolved:

1. The message will include the name of the file where the error occurred. Note the name of this file, as the data it contains may have been corrupted and the application using this file may encounter problems using this file's data.
2. If the file is not part of the SAN File System, refer to your system documentation for resolving file system errors.
3. If the file is part of the SAN File System:
   a. View the Event Log and resolve any errors that may be related to this problem.
   b. Suspect a communications problem either with the SAN or between the client and the Metadata server.

Related topics:

- "AIX client logging and tracing" on page 42
- "Windows 2000 client logging and tracing" on page 44
- "Troubleshooting the local network" on page 59
- Chapter 5, "Isolating problems with the SAN File System", on page 51
- "stfsdisk" on page 306
- "lsvol" on page 227
- "lsclient" on page 194

Troubleshooting client performance problems

Problem:

A client is either timing out or taking an unusually long time to access data.

Perform the following steps to resolve the problem:

1. Verify that the Metadata server that is managing the fileset being accessed is active. List the containers to view the server to which the fileset being accessed is attached, the Quota percentage and type, the attach point, and the directory.
   - From the Administrative command-line interface, run the **tanktool lscontainer** command.
   - From the SAN File System console, click **Manage Servers and Clients→Containers**.
2. Verify that the master Metadata server is active. Attempt to access master the Metadata server through either the SAN File System console or the Administrative command-line interface. If you cannot, suspect one of the following:
   - The master Metadata server is down. See Chapter 6, “Troubleshooting the cluster”, on page 55.

     Note: The Administrative agent will automatically attempt to restart a Metadata server that fails for any reason. Therefore, if you cannot access a Metadata server, you might want to wait a few minutes and try again just to ensure that it is not in the process of restarting.
   - There is an IP network problem between the client and the cluster. See Chapter 5, “Isolating problems with the SAN File System”, on page 51.

Related topics:

- lscontainer

Chapter 9. Troubleshooting the master console

The master console for the SAN File System is an IBM @server xSeries® 305 Type 8673 Model RA1 and provides support for Call Home and Remote Access.

Note: A Virtual Private Network (VPN) connection is required for Remote Access functionality. In addition, a remote display emulation package, such as Virtual Networking Computer (VNC), must be installed to access the SAN File System console or the RSA II Web interface.

The following software is used on the master console:

- Microsoft Windows 2000 Advanced Server edition

  Note: The optional Simple Network Management Protocol (SNMP) extensions should not be installed or, if they are installed, the SNMP Trap service must be disabled.
- IBM Director Server, version 4.1
- IBM Tivoli Bonus Pack for SAN Management
- Adobe Acrobat, version 5.0
- The PuTTY openssh package

For troubleshooting information, refer to the *IBM @server xSeries 305 Hardware Maintenance Manual and Troubleshooting Guide*. In addition, you can use the documentation that is available with each of the software packages loaded on the master console to troubleshoot software problems.

Related topics:

- "Service Alert" on page 16
- "Remote access" on page 17
- "Remotely accessing the master console" on page 20

Resolution procedures

You can use the procedures in this section to help you resolve problems with the master console. These procedures include:

- "Performing a total software recovery"
- "Recovering a hard disk drive" on page 84
- "Replacing fiber-channel cable and GBICs" on page 85

Performing a total software recovery

Steps:

The following steps explain how to perform a total software recovery of the master console using the set of recovery CDs:

1. Power OFF the master console.
2. Insert recovery CD 1.
3. Power ON the master console and follow the on-screen instructions.

4. Check each software package, updating to the latest level where required. Use the supplied CD or download the particular software package from the Web site.

Result:

The master console software is now reset to manufacturing default settings. Refer to the *SAN File System Planning and Installation guide* as well as the customer installation worksheets to update the master console to the current settings.

Related topics:

- "Master console" on page 16

Recovering a hard disk drive

Steps:

The hard disk drive in the master console is in the form of a mirrored pair to protect against the loss of the master console due to a disk failure. Perform the following steps to recover a master console hard disk drive:

1. Right-click the **My Computer** icon on your desktop and select **Manage**.
2. Select **Disk Management**. The hard drives are displayed in the right panel.
3. If the failing disk drive is displayed, right-click the main volume of the drive and select **Break Mirror**.

   Note: The mirror might have already been broken.
4. Shut down the master console and replace the failing disk drive using the procedures detailed in the xSeries 305 service documentation. Ensure that the new drive has its jumpers set the same as the drive that is being replaced. The new drive must be the same capacity or larger than the drive being replaced.

   Notes:

   a. It might not be obvious which of the two drives has failed. In this case, reboot with each drive connected in turn to isolate the failed drive.
   b. If the replacement drive has a boot record present, erase it prior to use.
   c. If the master console fails to boot because the Boot Record cannot be found, change the boot sequence in the BIOS to the other hard drive.
5. Disconnect the fibre-channel cables from the master console, making note of where they were connected.
6. Restart the master console.
7. Right-click the **My Computer** icon on your desktop and select **Manage**.
8. Select **Disk Management**. The hard drives are displayed in the right panel.
9. If a disk drive is displayed in the list marked "Missing," remove it by right-clicking the drive and selecting **Remove Disk**.
10. If the new disk drive has a "no entry sign" displayed on it, right-click it and select **Write Signature** to remove the "no entry sign."
11. Right-click the new disk drive and select **Upgrade to Dynamic Disk**.
12. Right-click the volume that you want to mirror and select **Add Mirror**. This step starts the Add Mirror Wizard.
13. Use the dialog boxes displayed to configure the second volume.
14. A dialogue box with reference to making changes to the boot.ini file is displayed. You can safely ignore this dialog box.

15. The status of both volumes, the existing drive, and the new drive will change to “Regenerating” and will, after a short period of time, start to show the percentage of regeneration completed. When the regeneration completes, the status changes to “Healthy.”
16. Reconnect the fibre-channel cables to the master console.

Related topics:

For more detailed information about installing, maintaining, and troubleshooting the master console, refer to *xSeries Type 8673 Installation Guide* and *xSeries 305 Type 8673 Hardware Maintenance Manual and Troubleshooting Guide*.

Replacing fiber-channel cable and GBICs

Steps:

Perform the following steps to replace a fibre-channel cable or Gigabit Interface Converter (GBIC):

1. Disconnect each end of the suspected failing fibre-channel cable.
2. Fit a replacement fibre-channel cable.
3. Check out the repair.
   - If the repair fixes the problem:
     a. Ensure that labels are fitted to each end of the new fibre-channel cable with the same information that was on the original fibre-channel cable.
     b. Follow customer procedures for the safe disposal of the original fibre-channel cable.
   - If the repair does not fix the problem, remove the new fibre-channel cable and reconnect the original fibre-channel cable.
4. Replace the GBICs on each end of the failing link, one at a time, and checking to see if the problem is resolved with each replacement. If a new GBIC does not resolve the problem, refit the original GBIC.

Related topics:

For more detailed information about installing, maintaining, and troubleshooting the master console, refer to *xSeries Type 8673 Installation Guide* and *xSeries 305 Type 8673 Hardware Maintenance Manual and Troubleshooting Guide*.

Chapter 10. Troubleshooting an engine

To perform engine troubleshooting procedures, you must have a monitor and keyboard, which should be provided by the customer. In most cases, you can perform the procedures in this section by remotely attaching through the master console to the RSA card that is installed in the engine.

If you suspect a problem with the RSA adapter that is installed in the engine, you will need to locally attach to the engine to resolve the problem.

You can verify that engine temperature, fans, and voltage are operating properly from either the Administrative command-line interface (**lsengine** or **statengine** command) or the SAN File System console (the Engine panel). You can also isolate problems to a specific engine.

Notes:

1. Before you begin troubleshooting an engine, it must first be taken offline from the SAN File System cluster. After diagnosing and repairing the engine, it must be brought back online. Verify with the customer that the engine is offline before attempting to troubleshoot or replace any hardware components.
2. Depending on the proximity of the master console to the engines in the SAN File System cluster, you may need to locally attach a customer-supplied keyboard, monitor, and mouse to the engine before performing the procedures in this section.

Related topics:

- "Accessing the Administrative server through a browser" on page 22
- "Accessing an engine through SSH" on page 23
- "Taking a Metadata server offline" on page 63

Performing general engine checkout

Prerequisites:

The Metadata server must be offline.

Steps:

Follow these steps to begin isolating hardware problems on an engine:

1. Power OFF the engine and all external devices.
2. Verify that all cables and power cords are properly connected.
3. If you are locally attached to the engine, set all display controls, such as contrast and brightness, to the middle position.
4. Power ON all external devices.
5. Power ON the engine.
6. Record any POST error messages that are displayed on the screen. If an error is displayed, see POST error messages for information about resolving that error.

   Note: If the engine detects no errors during POST, there will be one beep at the completion of the POST.

7. View the system error LED. If it is on, perform light path diagnostics.
8. View the system error log and resolve any error messages.
9. Run the diagnostic programs to isolate the component for which a problem is occurring.

Related topics:

- "POST error codes" on page 389
- "Light path diagnostic panel LEDs" on page 387
- "Running engine diagnostics" on page 104

Resolving hardware component problems

If you suspect problems with a specific component, use the following procedures to resolve those problems.

Note: If you have just installed new software or hardware and the engine is not working, complete these steps before attempting to use the problem resolution procedures:

1. Remove the software or hardware that you just added.
2. Run the diagnostic tests to verify that the engine is operating properly.
3. Reinstall the new software or hardware.

Related topics:

- "Performing general engine checkout" on page 87
- "Troubleshooting CD-ROM drive problems"
- "Troubleshooting diskette drive problems" on page 89
- "Troubleshooting the Ethernet controller" on page 90
- "Troubleshooting hard disk drive problems" on page 92
- "Troubleshooting keyboard, mouse, or pointing device problems" on page 93
- "Troubleshooting memory problems" on page 94
- "Troubleshooting microprocessor problems" on page 94
- "Troubleshooting monitor problems" on page 95
- "Troubleshooting power problems" on page 97
- "Troubleshooting USB port problems" on page 99

Troubleshooting CD-ROM drive problems

Problem:

Use this information to troubleshoot problems that you are having with your CD-ROM drive, such as:

- The engine does not recognize the CD-ROM drive.
- The CD-ROM drive is not working properly.
- The CD-ROM drive tray is not working.

Investigation steps:

Try the following actions to resolve the problem:

If the engine does not recognize the CD-ROM drive

Perform the following steps until you resolve the problem:

1. Verify that the IDE channel to which the CD-ROM drive is attached (primary or secondary) is enabled in the Configuration/Setup utility program.
2. Verify that all cables and jumpers are installed correctly according to the documentation that came with the CD-ROM drive.
3. Verify that the correct device driver is installed for the CD-ROM drive.
4. Run CD-ROM drive diagnostics and resolve any diagnostic error codes that result.
5. Replace the CD-ROM drive.

If the CD-ROM is not working properly

Perform the following steps in order until you resolve the problem:

1. Use the documentation that came with the CD-ROM drive to clean the CD-ROM.
2. Run CD-ROM drive diagnostics and resolve any diagnostic error codes that result.
3. Replace the CD-ROM drive.

If the CD-ROM tray is not working.

Perform the following steps in order until you resolve the problem:

1. Make sure the engine is powered ON.
2. Insert the end of a straightened paper clip into the manual tray-release opening.
3. Run CD-ROM drive diagnostics and resolve any diagnostic error codes that result.
4. Replace the CD-ROM drive.

Related topics:

- "POST error codes" on page 389
- "Running engine diagnostics" on page 104
- "Diagnostic error codes" on page 381
- "Replacing a CD-ROM drive" on page 119

Troubleshooting diskette drive problems

Problem:

Use this information to troubleshoot problems that you are having with your diskette drive.

Investigation steps:

Try the following actions to resolve the problem:

If the diskette drive LED remains lit or the engine bypasses the diskette drive

Perform the following steps until you resolve the problem:

Note: To prevent diskette read or write errors, make sure that the distance between monitors and diskette drives is at least 76 mm (3 in.).

1. Verify that a diskette is loaded in the diskette drive. Make sure that it is inserted correctly into the drive and that it is not damaged. (Try another diskette if you have one.)
2. Verify that the software is working properly.

3. Verify that the data cable (located at the rear of the diskette drive) is installed correctly.
4. Run the diskette drive diagnostics and resolve any diagnostic error codes that result.
5. Replace the data cable.
6. Replace the diskette drive.
7. Replace the system planar and shuttle assembly.

Related topics:

- "Running engine diagnostics" on page 104
- "Diagnostic error codes" on page 381
- "Replacing a diskette drive" on page 121

Troubleshooting the Ethernet controller

Problem:

Use this information to troubleshoot problems that you are having with the Ethernet controller, such as:

- The Ethernet controller cannot connect to the network.
- The engine stops running with loading device drivers.
- The Ethernet link status LED is not lit.
- The Ethernet transmit/receive activity LED is not lit.
- Data received are incorrect or sporadic.
- The Ethernet controller stopped working when another adapter was added to the engine.
- The Ethernet controller stopped working for no apparent reason.

The integrated Ethernet controller might display messages from certain device drivers. The latest information concerning these messages is available at the IBM Support Web site (http://www.ibm.com/pc/support).

Investigation steps:

Try the following actions to resolve the problem:

If the Ethernet controller cannot connect to the network

Perform the following steps until the problem is resolved:

1. Verify that the Ethernet cable is installed correctly. It must be securely attached at all connections.

   Note: If you set the Ethernet controller to operate at either 100 Mbps or 1000 Mbps, you must use Category 5 or higher cabling.
2. Replace the Ethernet cable with a different cable.
3. Determine whether the hub supports auto-negotiation. If it does not, configure the integrated Ethernet controller manually to match the speed and duplex mode of the hub.
4. Check the Ethernet controller LEDs, which are located on the rear of the engine, to see whether a problem exists with the connector, cable, or hub:

a. If the Ethernet activity LED is not lit, make sure the hub and network are operating and that the correct device drivers are installed.
b. If the Ethernet link status LED is not lit, there might be a defective connector, cable, or a problem with the hub.

5. Verify that you are using the correct device drivers, which were supplied with your engine.
6. Verify that there are no operating-system-specific problems related to the Ethernet controller. You can view the system error log, which is located at /var/log/messages to determine whether there are any problems.
7. Attempt to ping the IP address of the engine. If the ping is not successful, suspect an internal IP network configuration problem (a wrong IP address) or an external IP network configuration problem (the existence of a firewall).

If the engine stops running when loading device drivers

Perform the following steps until the problem is resolved:

1. Use the Configuration/Setup utility to verify that the interrupt (IRQ) setting assigned to the Ethernet controller is not also assigned to another device.

   Note: Interrupt sharing is allowed for PCI devices. However, some devices do not function well when they share an interrupt with a dissimilar PCI device.

   If the IRQ setting for the Ethernet controller is shared with another device, change the IRQ setting for either the Ethernet controller or the other device.
2. Verify that you are using the most recent device driver available from the Web.
3. Run the network diagnostics and resolve any diagnostic error codes that result.

If the Ethernet link status LED is not lit

Perform the following steps until the problem is resolved:

1. Verify that the hub is turned on.
2. Check all connections between the Ethernet controller and the hub.
3. Use another port on the hub.
4. Determine whether or not the hub supports auto-negotiation. If it does not, configure the integrated Ethernet controller manually to match the speed and duplex mode of the hub.
5. Run LED diagnostics and resolve any diagnostic error codes that result.
6. Reseat the adapter.
7. Replace the adapter.

If the Ethernet activity LED is not lit

Perform the following steps until the problem is resolved:

1. Attempt to send data from the engine to ensure the network is not idle.
2. Verify that you have loaded the network device drivers.
3. Run LED diagnostics and resolve any diagnostic error codes that result.

If the controller is receiving incorrect or sporadic data

Perform the following steps until the problem is resolved:

1. Verify that you are using Category 5 or higher cabling when operating the engine at 100 Mbps or 1000 Mbps.
2. Make sure that the cables are not physically close to noise-inducing sources, such as fluorescent lights.

If the Ethernet controller stopped working when another adapter was added to the engine

Perform the following steps until the problem is resolved:

1. Verify that the cable is connected to the Ethernet controller.
2. Make sure that the PCI system BIOS code is current.
3. Reseat the adapter.
4. Use the Configuration/Setup utility to verify that the interrupt (IRQ) setting assigned to the Ethernet controller is not also assigned to another device.

   Note: Interrupt sharing is allowed for PCI devices. However, some devices do not function well when they share an interrupt with a dissimilar PCI device.

   If the IRQ setting for the Ethernet controller is shared with another device, change the IRQ setting for either the Ethernet controller or the other device.
5. Replace the adapter.

If the Ethernet controller stopped working for no apparent reason

Perform the following steps until the problem is resolved:

1. Run diagnostics for the Ethernet controller and resolve any diagnostic error codes that result.
2. Try a different connector on the hub.
3. Reinstall the device drivers.
4. Reseat the adapter.
5. Replace the adapter.

Related topics:

- "Running engine diagnostics" on page 104
- "Diagnostic error codes" on page 381
- "Adding or replacing an adapter" on page 111

Troubleshooting hard disk drive problems

Problem:

Use this information to troubleshoot problems that you are having with a hard disk drive, such as:

- Some hard disk drives are not recognized during the hard disk drive diagnostic text (Fixed Disk test).
- The engine stops responding during the hard disk drive diagnostic text (Fixed Disk test).
- Data received are incorrect or sporadic.
- The Ethernet controller stopped working when another adapter was added to the engine.
- The Ethernet controller stopped working for no apparent reason.

Investigation steps:

Try the following actions to resolve the problem:

If some hard disk drives are not recognized during the hard disk drive diagnostic text (Fixed Disk test)

Perform the following steps in order until you resolve the problem:

1. Remove the first hard disk drive that was not recognized.
2. Run the Fixed Disk diagnostic test again.
3. If the remaining drives are recognized, replace the hard disk drive that you removed with a new hard disk drive.

If the engine stops responding during the hard disk drive diagnostic text (Fixed Disk test)

Perform the following steps in order until the problem is resolved:

1. Remove the hard disk drive being tested when the engine stopped responding.
2. Run the Fixed Disk diagnostic test again.
3. If the diagnostic test completes successfully, replace the hard disk drive that you removed with a new hard disk drive.

Related topics:

- "Running engine diagnostics" on page 104
- "Diagnostic error codes" on page 381
- "Replacing a hot-swap hard disk drive" on page 122

Troubleshooting keyboard, mouse, or pointing device problems

Problem:

Use this information to troubleshoot problems that you are having with your keyboard, mouse, or pointing device, such as:

- Some (or all) of the keys on the keyboard do not work.
- The mouse or pointing device does not work.

Investigation steps:

Try the following actions to resolve the problem:

If some (or all) of the keys on the keyboard do not work

Perform the following steps in order until you resolve the problem:

1. Verify that the keyboard cable is securely connected to the keyboard port (and not the mouse port) on the engine.
2. Verify that the engine and the monitor are turned on.
3. Replace the keyboard.
4. Replace the system planar and shuttle assembly.

If the mouse or pointing device does not work

Perform the following steps in order until you resolve the problem:

1. Verify that the mouse cable is securely connected to the mouse port (and not the keyboard port) on the engine.
2. Verify that the engine and the monitor are turned on.

3. Replace the mouse or pointing device.
4. Replace the system planar and shuttle assembly.

Related topics:

- "Replacing the system planar and shuttle" on page 142

Troubleshooting memory problems

Problem:

Use this information to troubleshoot problems that you are having with memory, such as the amount of system memory displayed is less than the amount of physical memory installed.

Investigation steps:

Try the following actions to resolve the problem:

If the amount of system memory displayed is less than the amount of physical memory installed

Perform the following steps until you resolve the problem:

1. Verify that:
   - The memory modules are seated properly.
   - You have installed the correct type of memory.
   - If you have changed the memory, you updated the memory configuration using the Configuration/Setup utility program.
   - All DIMM memory banks are enabled. The engine might have automatically disabled a DIMM memory bank because it detected a problem or a DIMM memory bank might have been manually disabled.
2. Check the POST error log for any 289 error message:
   a. If the DIMM was disabled by a system-management interrupt (SMI), replace the DIMM.
   b. If the DIMM was disabled manually or during POST:
      1) Start the Configuration/Setup utility program.
      2) Enable the DIMM.
      3) Save the configuration.
      4) Restart the engine.
3. Replace the DIMM.
4. Replace the system planar and shuttle assembly.

Related topics:

- "Replacing the system planar and shuttle" on page 142
- "Running engine diagnostics" on page 104
- "Engine error logs" on page 32

Troubleshooting microprocessor problems

Problem:

Use this information to troubleshoot problems that you are having with a microprocessor, such as the engine emits a continuous tone during POST.

Investigation steps:

Try the following actions to resolve the problem:

If the engine emits a continuous tone during POST

The startup (boot) microprocessor is not working properly. Perform the following steps in order until you resolve the problem:

1. Verify that the startup microprocessor is seated properly.
2. Replace the startup microprocessor.

Related topics:

- "Replacing a microprocessor" on page 131
- "Beep symptoms" on page 377

Troubleshooting monitor problems

Problem:

Use this information to troubleshoot problems that you are having with your monitor, such as:

- The monitor screen is blank.
- Only the cursor appears.
- The monitor works when you turn on the server but goes blank when you start some application programs.
- The screen is wavy, unreadable, rolling, distorted, or has screen jitters.
- The wrong characters appear on the screen.
- There is no video.

You should also refer to the documentation that comes with the monitor for adjustments that you can make and additional testing instructions.

Investigation steps:

Try the following actions to resolve the problem:

If the monitor screen is blank

Perform the following steps until you resolve the problem:

1. Verify that:
   - The engine power cord is plugged into the engine and a working electrical outlet.
   - The monitor power cord is plugged into a working electrical outlet.
   - The monitor cables are connected properly.
   - The monitor is turned on and the Brightness and Contrast controls are adjusted correctly.
   - If the servers are chained using IBM Cable Chaining Technology (C2T), verify these items:
     - The C2T chain cables are securely connected to the engines.
     - The C2T breakout cable is connected properly.
     - The selected engine is turned on.

Note: In some memory configurations, the 3–3–3 beep code might sound during POST, followed by a blank screen. If this occurs and the Boot Fail Count feature in the Start Options of the Configuration/Setup utility is set to Enabled (the default), you must restart the server three times to force the system BIOS to reset the CMOS values to the default configuration (memory connector or bank of connectors enabled).

2. Replace the monitor.
3. Replace the video adapter if installed.
4. Replace the system planar and shuttle assembly.

If only the cursor appears

Perform the following steps until you resolve the problem:

1. Run the diagnostics and resolve any diagnostic error codes that result.
2. Verify that the Power Supply LEDs indicate that the power supplies are working correctly.
3. Power off the engine.
4. Make sure the engine is cabled correctly.
5. Remove or disconnect the following devices (one at a time) until you find the failure (turn on the server and reconfigure each time):
   - Any external devices
   - Surge suppressor device (on the server)
   - Modem, printer, mouse, or non-IBM devices
   - Each adapter
   - Drives
   - Memory modules

   Note: Minimum operating requirements for the engine are:
   a. One power supply
   b. Power cage assembly
   c. System board
   d. One microprocessor and VRM
   e. Memory module (with a minimum of two 128 MB DIMMs)
6. Replace the power supply.
7. Replace the power cage assembly.
8. Replace the system planar and shuttle assembly.

If the monitor works when you turn on the server but goes blank when you start some application programs

Perform the following steps until you resolve the problem:

1. Verify that the application program is not setting a display mode higher than the capability of the monitor.
2. Verify that the primary monitor cable is connected to the C2T device breakout cable.
3. Verify that you have installed the device drives that are necessary for the application.
4. Replace the monitor.
5. Replace the video adapter.
6. Replace the system planar and shuttle assembly.

If the screen is wavy, unreadable, rolling, distorted, or has screen jitters

Perform the following steps until you resolve the problem:

1. Verify that the monitor self-tests show that the monitor is working properly.
   - If not, resolve the problems that occur during the self-tests.
   - If so, reconsider the location of the monitor. Magnetic fields around other devices (such as transformers, appliances, fluorescent lights, and other monitors) can cause screen jitter or wavy, unreadable, rolling, or distorted screen images.
     a. Turn off the monitor (moving a monitor while it is turned on may cause screen discoloration).
     b. Move the monitor at least 305 mm (12 in.) from the other device.
     c. Turn on the monitor.

     Notes:

     a. To prevent diskette drive read/write errors, make sure the distance between the monitor and diskette drives is at least 76 mm (3 in.).
     b. Non-IBM monitor cables might cause unpredictable problems.
     c. An enhanced monitor cable with additional shielding is available for the 9521 and 9527 monitors. For information about the enhanced monitor cable, contact your IBM reseller or IBM marketing representative.
2. Replace the monitor.
3. Replace the video adapter if installed.
4. Replace the system planar and shuttle assembly.

If the wrong characters appear on the screen

Perform the following steps until you resolve the problem:

1. If the wrong language is displayed, update the BIOS code with the correct language.
2. Replace the video adapter if installed.
3. Replace the system planar and shuttle assembly.

If there is no video

Perform the following steps until you resolve the problem:

1. Make sure the correct machine is selected, if applicable.
2. Make sure all cables are securely attached.

Related topics:

- "Replacing the system planar and shuttle" on page 142
- "Beep symptoms" on page 377
- "POST error codes" on page 389
- "Recovering BIOS code" on page 101

Troubleshooting power problems

Problem:

Use the information in this topic to troubleshoot problems that you are having with engine power or power supplies, such as:

- Neither the power switch nor the reset button work.
- The engine does not power on.

- Neither the AC nor the DC LED (located on the rear of the engine) are lit.
- The AC LED is lit, but the DC LED engine is not lit.
- The engine does not power off.

Investigation steps:

Try the following actions to resolve the problem:

If neither the power switch nor the reset button work

Perform the following steps until you resolve the problem:

1. Reseat the connector.
2. Replace the power switch card.
3. Replace the system planar and shuttle assembly.

If the engine does not power on

Perform the following steps until you resolve the problem:

1. Verify that the power cables are properly connected to the engine.
2. Verify that the electrical outlet is functioning properly.
3. Verify that the type of memory installed is correct.
4. If you have just installed an option into the engine, remove it and restart the engine. If the engine powers on, you might have installed more options than the power supply supports.
5. Override the front panel power button by turning on switch 3 of SW1:
   a. Disconnect engine power cords.
   b. Reconnect engine power cords.

   If the engine powers on, replace the operator information panel. If the engine does not power on, replace the system planar and shuttle assembly.

If neither the AC nor the DC LED (located on the rear of the engine) is lit

The engine has no power or there is an AC power problem. Perform the following steps until you resolve the problem:

1. Verify that the electrical cord is properly connected to the engine.
2. Verify that the electrical outlet is functioning properly.
3. Disconnect the ribbon cable from connector J25 on the system board. If the AC power LED comes on, see "Isolating undetermined problems" on page 101.
4. Replace the power supply.

If the AC LED is lit, but the DC LED engine is not lit

The engine is in standby mode or there is a DC power problem. Perform the following steps until you resolve the problem:

1. If the AC LED is flashing, press the power-control button to power on the engine.
2. Check system board cable connectors J4 and J10. Override the front panel pushbutton by turning on switch 3 of SW1. If the DC LED is lit:
   a. Press Ctrl+Alt+Delete.
   b. Watch the screen for any POST errors.
   c. Check the system error log.
   d. If the engine starts with no errors:
      1) Replace the power switch assembly.

2) Replace the system planar and shuttle assembly.

3. Remove the adapters and disconnect cables and power connectors to all internal and external devices.
4. Power ON the engine.
5. If the DC LED is lit, replace the adapters and devices, one at a time, until you isolate the problem.
6. Replace the power supply.
7. Replace the power cage assembly.
8. Replace the system planar and shuttle assembly.

If the engine does not power off

Perform the following steps until you resolve the problem:

1. Verify whether you are using an ACPI or a non-ACPI operating system. If you are using a non-ACPI operating system:
   a. Press Ctrl + Alt + Delete.
   b. Power OFF the engine by pressing and holding the power-control button for four seconds.
   c. If the engine fails during BIOS POST and the power-control button does not work, remove the AC power cord.
2. If the problem persists or if you are using an ACPI-aware operating system, suspect the system planar and shuttle assembly.

Related topics:

- "Running engine diagnostics" on page 104
- "Diagnostic error codes" on page 381
- "Troubleshooting power-on problems" on page 100
- "Isolating undetermined problems" on page 101

Troubleshooting USB port problems

Problem:

Use this information to troubleshoot problems that you are having with the Universal Serial Bus (USB) port, such as a USB device does not operate properly.

Investigation steps:

Try the following actions to resolve the problem:

If a USB device does not operate properly

Perform the following steps in order until you resolve the problem:

1. Verify that you are not trying to use a USB device during POST if you have a standard (non-USB) keyboard attached to the keyboard port. If you have a standard keyboard attached to the keyboard port, the USB port is disabled and no USB device will work during POST.
2. Verify that the correct USB device driver is installed.
3. Verify that your operating system supports USB devices.

Related topics:

- "Universal Serial Bus version 1.1 port" on page 375
- "Input and output ports" on page 372

Troubleshooting power-on problems

Problem:

Use this information to troubleshoot problems that you are having when you power on an engine, such as:

- Nothing happens when you power on the engine.
- POST completes successfully, but you do not hear any beeps.
- POST does not complete successfully.

When you power on an engine, it performs a series of tests, called the power-on self test (POST), that check the operation of each of the components as well as some of the options that are installed in the engine. If POST completes without detecting any problems, a single beep sounds and the first panel of your application appears. If POST detects a problem, more than one beep sounds, and an error message appears on the screen.

Investigation steps:

Try the following actions to resolve the problem:

If nothing happens when you power on the engine

Perform the following steps until you resolve the problem:

1. Verify that the power cables are properly connected to the engine.
2. Verify that the electrical outlet is functioning properly.
3. Verify that the type of memory installed is correct.
4. If you have just installed an option into the engine, remove it and restart the engine. If the engine powers on, you might have installed more options than the power supply supports.
5. Override the front panel power button by turning on switch 3 of SW1:
   a. Disconnect engine power cords.
   b. Reconnect engine power cords.

   If the engine powers on, replace the operator information panel. If the engine does not power on, replace the system planar and shuttle assembly.

If POST completes successfully, but you do not hear any beeps

Perform the following steps until you resolve the problem:

1. Run the Configuration/Setup utility and verify that Start Options Power-On Status is enabled.
2. Replace the system planar and shuttle assembly.

If POST does not complete successfully

Use either the beeps that you hear or the POST error codes that you see to resolve the problem.

Related topics:

- "Beep symptoms" on page 377
- "POST error codes" on page 389
- "Troubleshooting power problems" on page 97
- "Isolating undetermined problems" on page 101
- "Replacing the system planar and shuttle" on page 142

Isolating undetermined problems

Prerequisites: Use the information in this section if the diagnostic tests did not identify the failure, the device list is incorrect, or the system is inoperative.

Note: Damaged data in CMOS or BIOS can cause undetermined problems.

Steps:

Perform the following steps to isolate undetermined problems with an engine:

1. Verify that the power supply LEDs indicate that the power supplies are working correctly.
2. Power OFF the engine.
3. Make sure the engine is cabled correctly.
4. Remove or disconnect the following devices (one at a time) until you find the failure (turn on the server and reconfigure each time):
   - Any external devices
   - Surge suppressor device (on the server)
   - Modem, printer, mouse, or non-IBM devices
   - Each adapter
   - Drives
   - Memory modules

   Note: Minimum operating requirements for the engine are:
   a. One power supply
   b. Power cage assembly
   c. System board
   d. One microprocessor and VRM
   e. Memory module (with a minimum of two 128 MB DIMMs)
5. If the problem remains, suspect the following FRUs in the order listed:
   - Power supply
   - Power cage assembly
   - System planar and shuttle assembly

Result:

If the problem goes away when you remove an adapter from the system and replacing that adapter does not correct the problem, suspect the system planar and shuttle assembly.

If you suspect a networking problem and all the system tests pass, suspect a network cabling problem external to the system.

Related topics:

- "Replacing a hot-swap power supply" on page 125
- "Replacing the system planar and shuttle" on page 142
- "Replacing the power-supply cage" on page 137
- "Diagnostic error codes" on page 381

Recovering BIOS code

Context:

If the BIOS code has become damaged, such as from a power failure during a flash update, you can recover the BIOS code using the boot block jumper and a BIOS flash diskette.

Note: The flash memory of your engine contains a protected area that cannot be overwritten. The recovery boot block is a section of code in this protected area that enables the engine to start up and read a flash diskette. The flash utility recovers the system BIOS code from the BIOS recovery files on the diskette.

Steps:

Perform the following steps to recover BIOS code:

1. Obtain a BIOS flash diskette. You can obtain a BIOS flash diskette from one of the following sources:
   - The *ServerGuide Setup and Installation* CD.
   - The TotalStorage 4146 support Web site at: www.ibm.com/storage/support/4146.
   - Your IBM service representative.
2. Power off the engine.
3. Disconnect all power cords and external cables.
4. Remove the engine cover.

5. Locate the flash boot recovery jumper block (J34) on the system board.

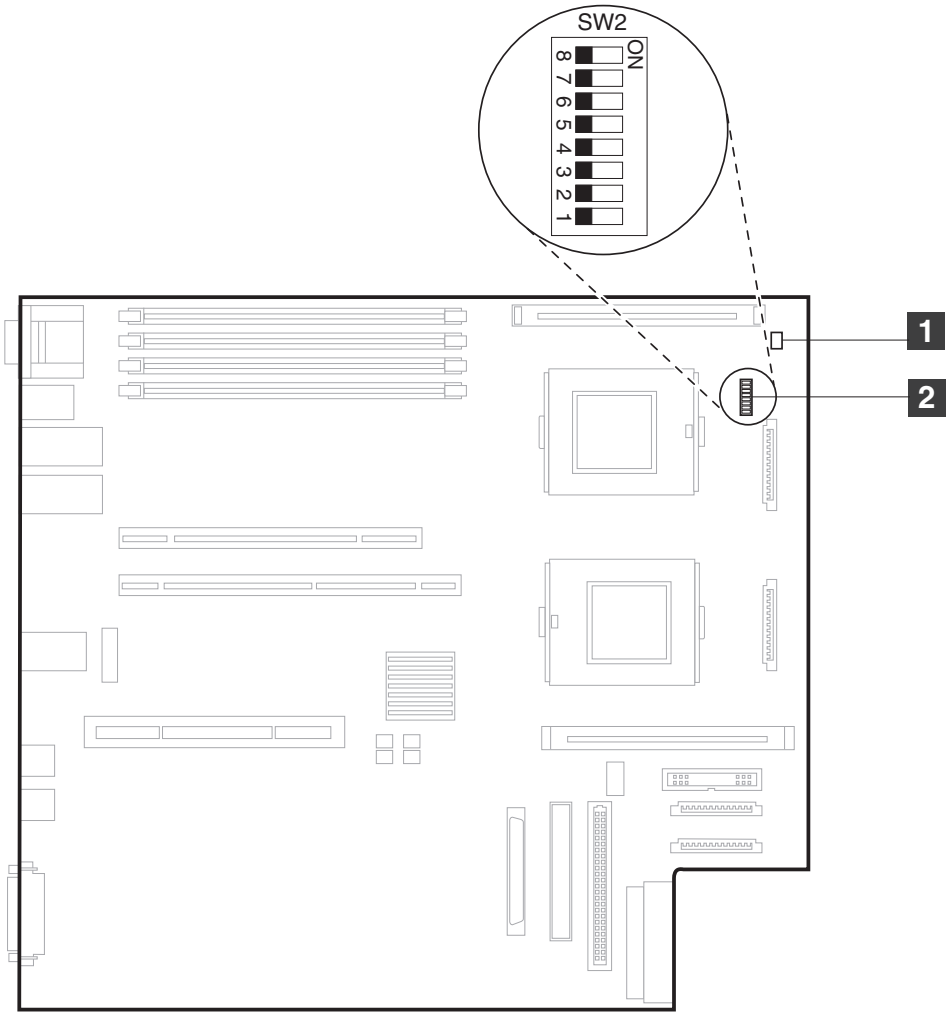

1 Flash boot book recovery jumper (J34)

2 System board switch block (SW2)

6. Move the jumper from pins 1 and 2 to pins 2 and 3 to enable the BIOS recovery mode.
7. Insert the BIOS flash diskette into the diskette drive.
8. Replace the engine cover, and reconnect all power cords.
9. Power ON the engine. It begins POST.
10. Select **1 - Update POST/BIOS** from the menu that contains various flash update options.
11. When prompted as to whether you want to save the current code to a diskette, type `N`.
12. When prompted to choose a language, select a language (by entering a number from 0 to 7) and press Enter to accept your choice.
13. Remove the BIOS flash diskette from the diskette drive.
14. Power OFF the engine.
15. Disconnect all power cords and external cables.
16. Remove the engine cover.

17. Remove the jumper from the Flash boot block recovery jumper block, or move it to pins 1 and 2 to return to normal startup mode.
18. Reconnect all external cables and power cords, and turn on the peripheral devices.
19. Replace the engine cover.
20. Power ON the server. The system should start up normally.

Related topics:

- "Removing the cover and bezel" on page 110

Running engine diagnostics

Steps:

To run the diagnostic tests for an engine:

1. Power on the server and watch the screen.

   Note: To run the diagnostic tests, you must use the highest level password available for the engine when you power it on. For example, if an administrative password is set for the engine, you must use it when you power on the engine instead of using the power-on password.
2. When the message `F2 for Diagnostics` appears, press F2.
3. Type the appropriate password and press Enter.
4. Select **Extended** from the top of the panel.
5. When the Diagnostic Programs panel appears, select the test that you want to run from the list and follow the instructions on the screen.

Notes:

1. Press F1 while running the diagnostic programs to obtain help information. Press F1 within a help screen to obtain online documentation. To exit from the help information and return to where you were, press Esc.
2. If the engine stops during testing and you cannot continue, restart the engine and try to run the diagnostic programs again. If the problem persists, replace the component that was being tested when the server stopped.
3. The keyboard and mouse (pointing device) tests assume that a keyboard and mouse are attached to the engine. You can test a USB keyboard using the regular keyboard test. In addition, you can test a USB mouse using the regular mouse test.
4. If you run the diagnostic programs with either no mouse or a USB mouse attached to the engine, you will not be able to navigate between test categories using the **Next Cat** and **Prev Cat** buttons. All other functions provided by mouse-selectable buttons are also available by using the function keys.
5. You can run the USB interface test only if there are no USB devices attached to the engine.
6. Select **Hardware Info** from the top of the panel to view server configuration information, such as system configuration, memory contents, interrupt request (IRQ) use, direct memory access (DMA) use, and device drivers.

Result:

When the tests are complete, the results are stored in the test log. You can view the test log only while you are in the diagnostic programs. When you exit the diagnostic programs, the test log is cleared unless you save it.

To view the test log:

1. Select **Utility** from the top of the panel.
2. Select **View Test Log**.

To save the test log so that you can view it later, click **Save Log** on the Diagnostics Program. Then, specify a location and name for the saved log file. To save the test log to a diskette, you must use a diskette that you have formatted yourself. You cannot save the test log to a preformatted diskette. The diskette can contain other data as long as there is enough room for the test log as well.

Related topics:

• "Diagnostic error codes" on page 381

Running OS Present diagnostics

Steps:

To run the OS Present diagnostic tests for an engine:

1. From a command line, run the **diags** command.
2. Follow the directions on the initial menu to adjust the size of the menus to your display.
3. From the initial menu, enter **1** to view the list of devices.

   Note: To obtain help on using the OS Present diagnostics, enter **h**.
4. The Device List menu displays all installed devices (even if those devices are currently missing). Enter a number that corresponds to a details link to display the Device Status and Control menu for that device.
5. The Device Status and Control menu shows the operational state, test state, and fault state for the device. The Operational state must be Service to run the test for a device.
6. Change the operational state of the engine to Service, if necessary:
   a. Enter **8** for reboot or shutdown .
   b. On the Reboot/Shutdown Confirmation menu, press **2** to reboot the engine in the Service state.
   c. From the command line, run the **diags** command
   d. Enter **1** to view the list of devices
   e. Enter the number that corresponds to the details link to display the Device Status and Control menu for the device.
7. From the Device Status and Control menu, choose the task to be performed:
   - Run test one time
   - Loop test until an error occurs
   - Loop test until manually stopped – log all errors
   - Stop looping test (if a looping test is currently being run)
   - Flash the blue "identity" LED on this engine for 2 minutes
   - View Diagnostic Error Log for this engine
   - View Remote Supervisor Adapter's event log

- View Service History log
- Reboot or shutdown

Note: When running the test for the Fiber Channel Adapter, you have the option of testing with a wrap plug (shipped with the SANFS system) connected to the adapter's port (or ports). In addition, a coupler (also shipped with the SANFS system) can be used to place the wrap plug at the end of the fiber cable and test the complete path to the switch.

8. Reboot the engine after running tests to return to normal operation.

Related topics:

- "OS Present diagnostics" on page 33
- "The diags command" on page 33

Logging on to the RSA II

Steps:

Perform the following steps to log on to the RSA II for an engine:

1. From the master console, point a Web browser to the IP address or host name of the RSA II to which you want to connect.
2. Enter the user name and password for the RSA II. Then click **OK**.

   Note: The RSA II is set initially with a user name of USERID and password of PASSW0RD (with a zero, not the letter O). This user has read/write access. The user name you use must have read/write access to the RSA II if you intend to redirect the engine console to your computer.
3. Select a timeout value from the RSA II Welcome page and click **Continue**.
4. The RSA II System Status page is displayed. Choose a function on the left navigation pane to continue.

Related topics:

- "RSA II Web interface" on page 35

Starting the Configuration/Setup Utility

Steps:

The Configuration/Setup Utility program is a menu-driven utility that is part of the BIOS code that comes with the Model 1RX. You can use it to perform these functions:

- Configure serial port assignments.
- Change the drive startup sequence.
- Enable USB keyboard and mouse support.
- Resolve configuration conflicts.
- Set the date and time.
- Set passwords.

Perform the following steps to start the Configuration/Setup Utility program:

1. Power ON the engine and watch the monitor screen.

2. When the message `Press F1 for Configuration/Setup appears`, press F1.

   Note: To run the Configuration/Setup Utility, you must use the highest-level password available for the engine when you power it on. For example, if an administrative password is set for the engine, you must use it when you power on the engine instead of using the power-on password.

   If you do not use the administrative password, you will have access to only a limited set of Configuration/Setup Utility functions.
3. Follow the instructions that appear on the screen.

Chapter 11. Adding and replacing engine components

Use the following procedures to add or replace engine components. Before you add or replace components in an engine, you must first prepare the engine by removing it from the SAN File System cluster.

Note: Before you begin the troubleshooting an engine, it must first be taken offline from the SAN File System cluster. After diagnosing and repairing the engine, it must be brought back online. Taking an engine offline from the SAN File System cluster requires several commands that can disrupt the operation of the SAN File System. Only qualified software support personnel should issue those commands.

- "Removing the cover and bezel" on page 110
- "Adding or replacing an adapter" on page 111
- "Replacing the battery" on page 117
- "Replacing a CD-ROM drive" on page 119
- "Replacing the hard disk backplane" on page 120
- "Replacing a diskette drive" on page 121
- "Replacing a hot-swap hard disk drive" on page 122
- "Replacing a hot-swap fan" on page 123
- "Replacing a hot-swap power supply" on page 125
- "Replacing a memory module" on page 127
- "Replacing a microprocessor" on page 131
- "Replacing the diagnostics/operator panel card" on page 136
- "Replacing the power-supply cage" on page 137
- "Replacing an RSA II adapter" on page 139
- "Replacing the system planar and shuttle" on page 142
- "Replacing the cover and bezel" on page 144

Removing the cover and bezel

Attention: For correct cooling and airflow, replace the cover before powering on the engine. Operating the engine for extended periods of time (over 30 minutes) with the cover removed might damage components due to incorrect cooling.

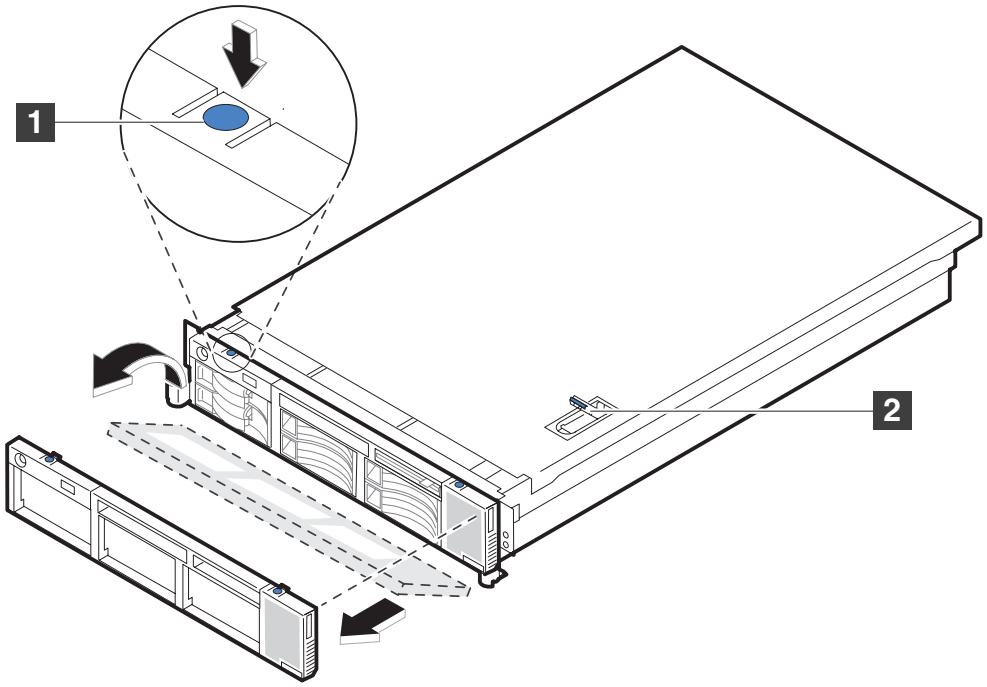

1 Bezel-release tab
2 Cover-release latch

Steps:

Perform the following steps to remove the top cover:

1. Review the safety information.
2. If you are planning to install or remove a microprocessor, memory-module, PCI adapter, or battery, power OFF the engine and any peripheral devices, and disconnect all external cables and power cords.
3. Release the left- and right-side latches and pull the engine out of the rack cabinet until both slide rails lock.

   Note: You can reach the cables on the back of the engine when the engine is in the locked position.
4. Lift the cover-release latch. Lift the cover off of the engine and set the cover aside.

Perform the following steps to remove the bezel:

1. Press in on the tabs on the top of the bezel.
2. Rotate the bezel away from front of the engine.
3. Store the bezel in a safe place.

Note: You do not need to remove the bezel to install or replace hot-swap hard disk drives in the Model 1RX.

Related topics:

- "Before you begin replacing components" on page 426
- "Handling static-sensitive devices" on page 424
- "Replacing the cover and bezel" on page 144
- "Safety information" on page 418

Adding or replacing an adapter

You can install up to five peripheral component interconnect (PCI) adapters in the connectors on the system planar and PCI riser card. The standard and optional adapters must be placed in the following PCI slots:

Slot 1 (open)
Slot 2 Standard 2-port Fibre Channel Adapter
Slot 3 Optional 1-port Fiber Ethernet Adapter
Slot 4 Optional 1-port Fiber Ethernet Adapter
Slot 5 Standard Remote Supervisor Adapter II (RSA II)

The following illustrations show the location of the PCI and PCI-X adapter expansion slots on the system planar and PCI riser card.

System planar:

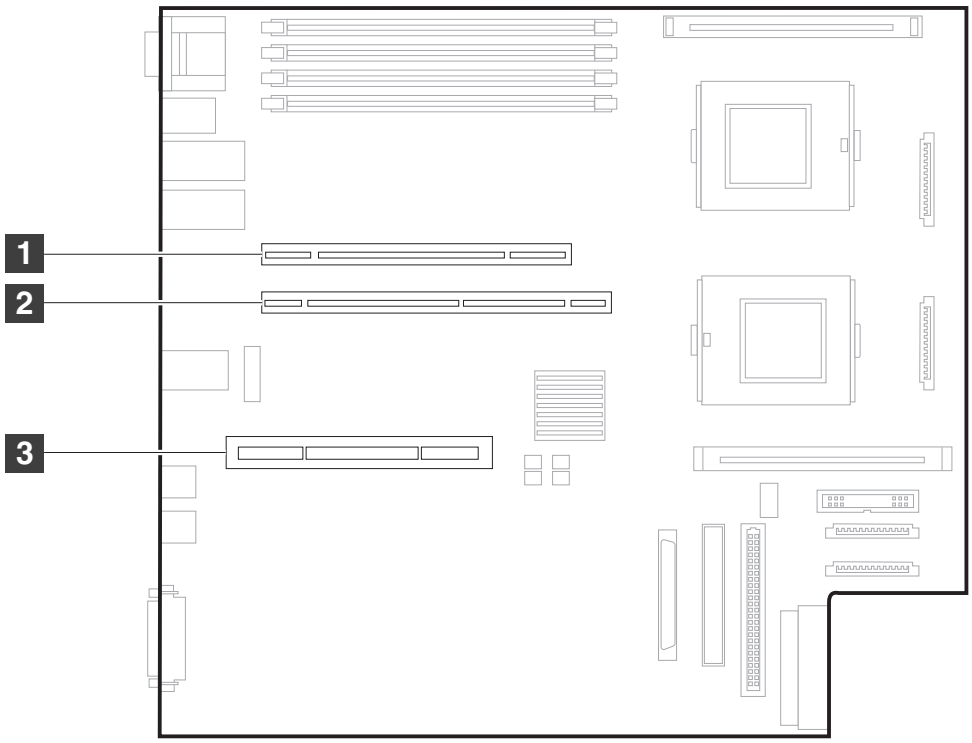

1 PCI-X slot 1 64–bit 3.3V 100 MHz (J14)
2 PCI-X slot 2 64–bit 3.3V 100 MHz (J13)
3 PCI riser card (J16)

PCI riser card:

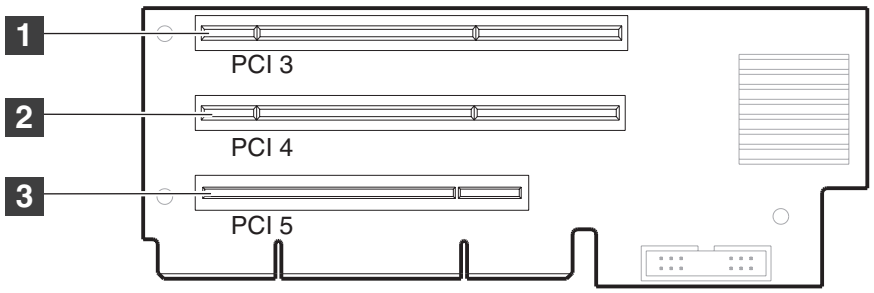

1 PCI-X slot 3 64–bit 3.3V 133 MHz (PCI3)
2 PCI-X slot 4 64–bit 3.3V 133 MHz (PCI4)
3 PCI-X slot 5 32–bit 5V 33 MHz (PCI5)

Adapter considerations:

Before you install an adapter, review the following information:

- Read the documentation that comes with your operating system.
- Locate the documentation that comes with the adapter and follow those instructions to supplement the following instructions. If you need to change the switch or jumper settings on your adapter, follow the instructions that come with the adapter.
- You can install only low-profile adapters in slots 1 and 2.
- You can install standard full-length adapters, with the adapter guides removed from the adapters, in slots 3 and 4.
- You can install only a half-length adapter in slot 5.
- The Model 1RX supports 5.0 V and universal PCI adapters on the 32-bit PCI slot 5 and supports only 3.3 V and universal PCI adapters in the 64-bit PCI/PCI-X slot 3 and 4.
- The PCI bus configuration is:
  - Non-hot-plug, 64-bit PCI-X slots 1 through 2 (PCI bus A, 100 MHz)
  - Non-hot-plug, 64-bit PCI-X slots 3 through 4 (PCI bus B, 133 MHz)
  - Non-hot-plug, 32-bit PCI slot 5 (PCI bus C, 33 MHz)
- The RSA II card comes preinstalled in PCI slot 5. It can be installed only in PCI slot 5. Use the ribbon cable that comes with the adapter to connect the adapter to the system management connector on the PCI riser-card.
- The system scans PCI and PCI-X slots 1 through 5 to assign system resources; then the system starts (boots) the system devices in the following order, if you have not changed the default boot precedence: integrated Ethernet controller, integrated SCSI controller, and then PCI and PCI-X slots 5, 3, 4, 1, and 2.

  Notes:

  1. To change the boot precedence for PCI and PCI-X devices, start the Configuration/Setup Utility program and select **Start Options** from the main menu.

Prerequisites:

A Phillips screwdriver may be required if you are replacing an adapter on the PCI-riser card assembly.

Steps:

The following illustration shows how to install a PCI or PCI-X adapter.

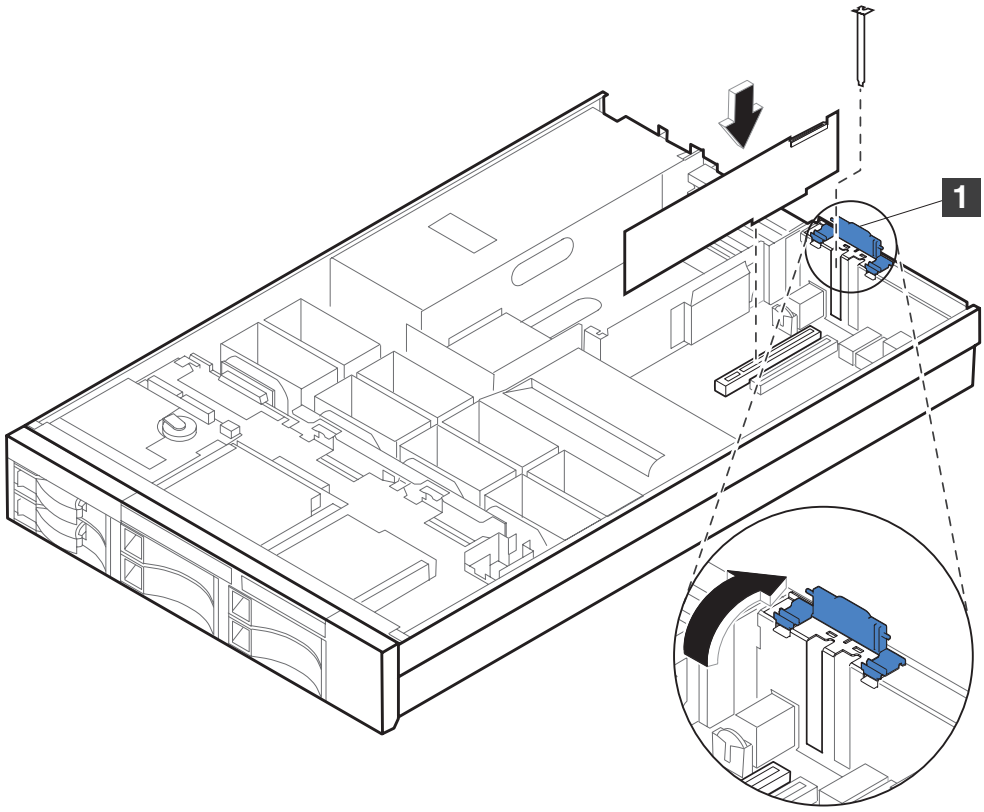

1 Adapter retainer

Attention: Expansion-slot covers must be installed on all empty slots. This maintains the electronic emissions characteristics of the system and ensures correct cooling of system components.

Perform the following steps to install or replace a PCI or PCI-X adapter:

1. Review the safety information and information about handling static-sensitive devices.
2. Power OFF the engine and any peripheral devices, and disconnect all external cables and power cords.
3. Remove the cover.
4. Determine which expansion slot you want to use for the adapter.

5. If you are replacing an adapter in PCI slot 3, 4, or 5, remove the PCI riser-card assembly:

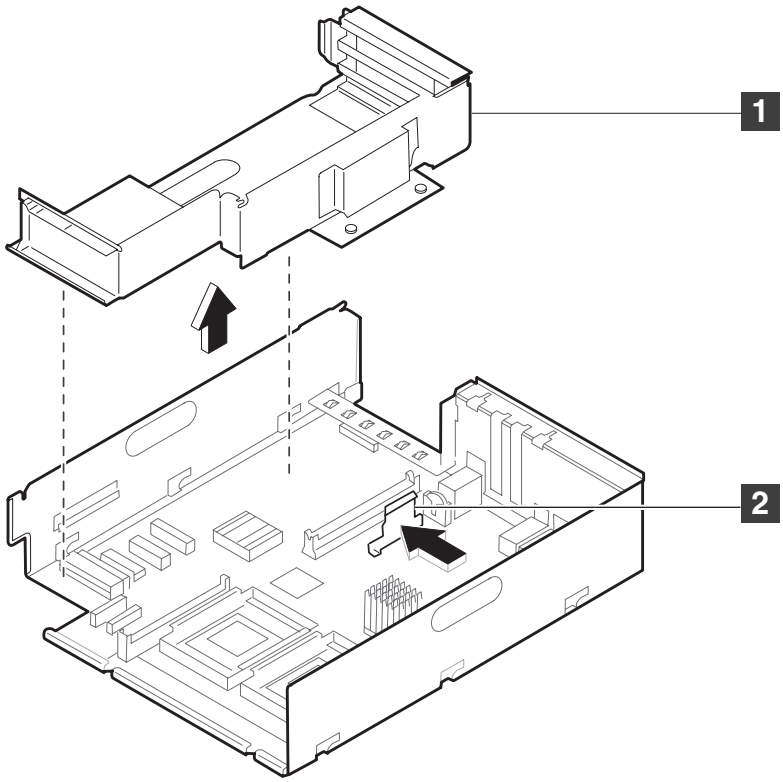

1 PCI-riser card assembly
2 Retention latch

   a. If the PCI riser-card assembly has a latch, release the latch. Otherwise, loosen the two captive screws to release the PCI riser-card assembly.
   b. Lift the assembly from the engine.
   c. Place the assembly on a flat, level surface.
6. Remove the defective adapter (or the adapter blank):
   a. Lift up the blue adapter-retainer.
   b. Carefully lift out the adapter.
   c. If an adapter cable is present, disconnect the adapter cable by squeezing the retention latches and pulling away from the adapter.
7. Refer to the documentation that comes with your adapter for any cabling instructions. It might be easier for you to route cables before you install the adapter.
8. Touch the static-protective package containing the adapter to an unpainted metal surface, such as the expansion-slot cover on the engine. This discharges any static electricity from the package and from your body.

   Attention: Avoid touching the components and gold-edge connectors on the adapter.
9. Remove the adapter from the static-protective package, holding the adapter by the edges.
10. Replace the adapter:

a. If you are installing a full-length adapter, remove the adapter guide (if any) from the end of the adapter.

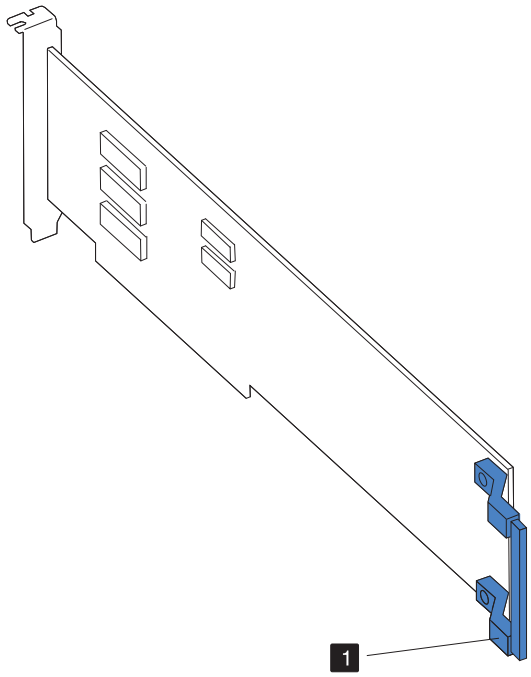

1 Adapter guide

b. Reconnect any necessary cables to the adapter. The cable connector is keyed for correct insertion. (Refer to the documentation that comes with your adapter for cabling instructions.)

Attention:
When you route the cables, do not block the ventilated space around any of the fans. In addition, route any cables that are plugged into the PCI adapter under adjacent adapters so they are not pinched between the top edge of the adapter and the shuttle sidewall.

c. Carefully grasp the adapter by its top edge or upper corners and align it with the expansion slot on the system planar.

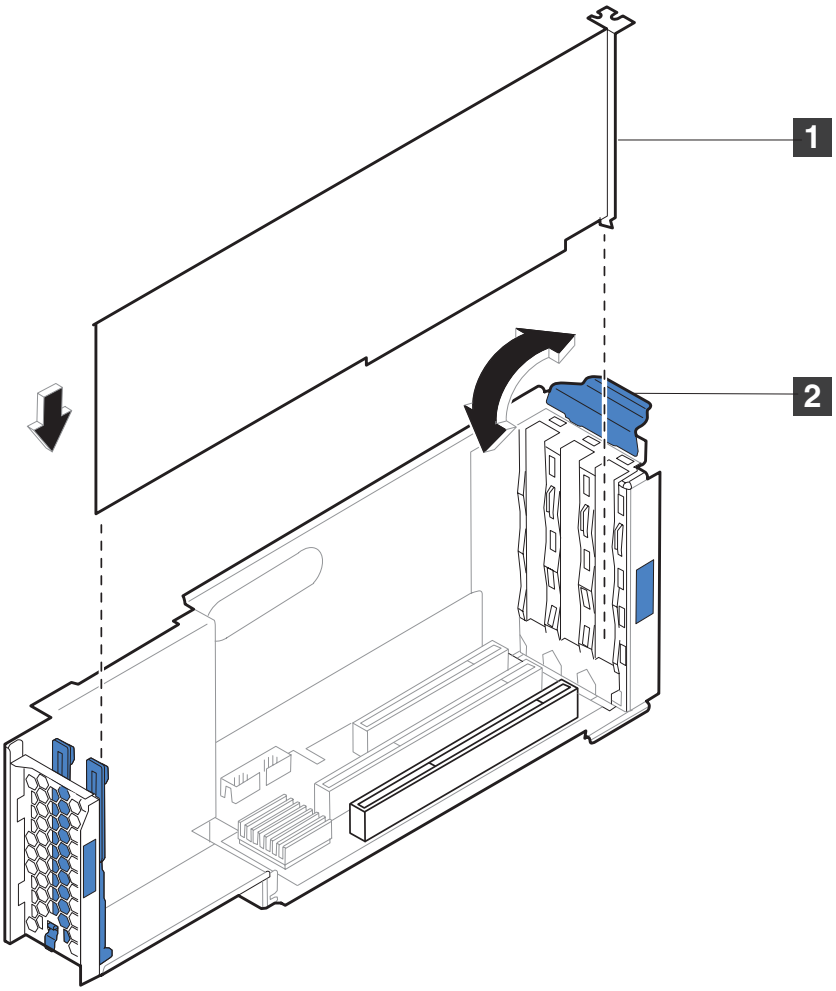

1 PCI adapter
2 Adapter retainer

d. Press the adapter ***firmly*** into the expansion slot and close the adapter retainer.

Attention: When you install an adapter, be sure that it is completely and correctly seated in the system planar connector before you apply power. Incomplete insertion might cause damage to the system planar or the adapter.

11. If you removed the PCI riser-card assembly to install the adapter,
    a. Align the metal bracket on the side of the PCI riser-card assembly with the blue tab next to the expansion slot so that the tab sits between the metal bracket and the PCI riser-card assembly.
    b. Press the PCI riser-card assembly ***firmly*** into the connector.
    c. Turn the retention latch and snap it into place to hold the PCI riser-card, or if captive screws are present, tighten the two captive screws.

    Note: Make sure that the PCI riser-card assembly is correctly seated on the system-planar connector and that the cables under the PCI riser-card assembly are not pinched.

12. Replace the cover, reconnect external cables and power cords, and power on the engine.

Related topics:

- "Before you begin replacing components" on page 426
- "Handling static-sensitive devices" on page 424
- "Replacing the cover and bezel" on page 144
- "Removing the cover and bezel" on page 110
- "Safety information" on page 418
- "Starting the Configuration/Setup Utility" on page 106

Replacing the battery

IBM has designed this product with your safety in mind. The lithium battery must be handled correctly to avoid possible danger. If you replace the battery, you must adhere to the following instructions.

CAUTION:
When replacing the lithium battery, use only IBM Part Number 33F8354 or an equivalent type battery recommended by the manufacturer. If your system has a module containing a lithium battery, replace it only with the same module type made by the same manufacturer. The battery contains lithium and can explode if not properly used, handled, or disposed of. (C23)

Do not:

- Throw or immerse into water
- Heat to more than 100°C (212°F)
- Repair or disassemble

Dispose of the battery as required by local ordinances or regulations.

Note: For translations of this safety notice, refer to *Translated Safety Notices*, GA67-0043.

Note: In the U.S., call 1-800-IBM-4333 for information about battery disposal.

If you replace the original lithium battery with a heavy-metal battery or a battery with heavy-metal components, be aware of the following environmental considerations. Batteries and accumulators that contain heavy metals must not be disposed of with normal domestic waste. They will be taken back free of charge by the manufacturer, distributor, or representative, to be recycled or disposed of in a correct manner. To order replacement batteries, call your IBM reseller or IBM marketing representative.

Note: After you replace the battery, you must reconfigure the engine and reset the system date and time.

Steps:

Perform the following steps to replace the battery:

1. Review the safety information.
2. Review any special handling and installation instructions supplied with the battery.
3. Power OFF the engine and any peripheral devices, and disconnect all external cables and power cords.
4. Remove the cover.
5. Remove the adapters from PCI slots 1 through 2.
6. Disconnect any internal cables, as necessary.
7. Locate the battery (connector BH1) on the system planar.

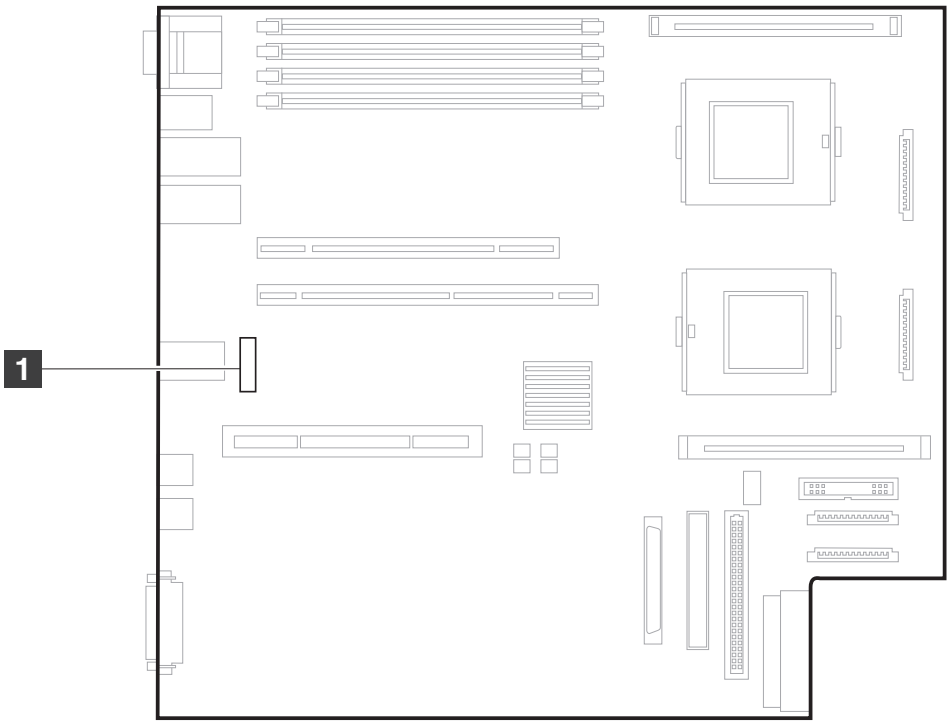

1 Battery (BH1)

8. Remove the battery:
   a. Slide your fingernail under the clip at the top of the battery and press the top of the battery clip away from the battery.
   b. Lift and remove the battery from the socket.

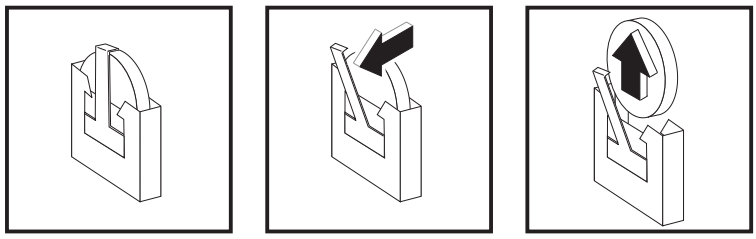

9. Insert the new battery:

   Note: Ensure that the polarity of the battery is correct. The positive (+) side must face the battery clip.

   a. Use one finger to press the top of the battery clip away from the battery.

b. Press the battery into the socket until it clicks into place. Make sure the battery clip holds the battery securely.

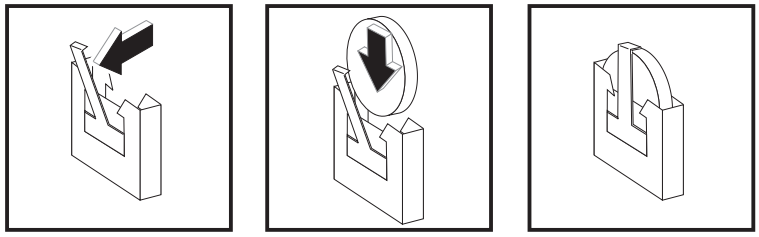

10. Reinstall the adapters that you removed.
11. Connect the internal cables that you disconnected.
12. Connect all external cables and all power cords.
13. Install the cover.
14. Start the Configuration/Setup Utility program and set the following parameters:
    - Set the system date and time.
    - Set the power-on password.
    - Reconfigure the system.

Related topics:

- "Before you begin replacing components" on page 426
- "Replacing the cover and bezel" on page 144
- "Removing the cover and bezel" on page 110
- "Adding or replacing an adapter" on page 111
- "Safety information" on page 418
- "Starting the Configuration/Setup Utility" on page 106

Replacing a CD-ROM drive

This topic describes how to replace the CD-ROM drive.

Note: The appliance's electromagnetic interference (EMI) integrity and cooling are both protected by having all drive bays covered or occupied.

CAUTION:
When laser products (such as CD-ROMs, fiber optic devices, or transmitters) are installed, note the following:

- **Do not remove the covers. Removing the covers of the laser product could result in exposure to hazardous laser radiation. There are no serviceable parts inside the device.**
- **Use of controls or adjustments or performance of procedures other than those specified herein might result in hazardous radiation exposure.**

DANGER

Some laser products contain an embedded Class 3A or Class 3B laser diode. Note the following:

Laser radiation when open. Do not stare into the beam, do not view directly with optical instruments, and avoid direct exposure to the beam. (C11)

Note: For translations of this safety notice, refer to *Translated Safety Notices*, GA67-0043.

Steps:

Perform the following steps to replace a CD-ROM drive:

1. Review the safety information and information about handling static-sensitive devices.
2. Power OFF the engine and peripheral devices. Disconnect all power cords and external cables.
3. Remove the cover.
4. Remove the defective CD-ROM drive:
   a. Remove the wire spring from the right side of the drive.
   b. Slide the drive out of the bay.
5. Install the new CD-ROM drive:
   a. Set the jumpers to the slave setting according to the documentation that comes with the drive.
   b. Align the rails on the drive with the guide rails in the drive bay.
   c. Gently push the drive into the bay until it connects to the card assembly.
6. Replace the cover and bezel, reconnect all power cords and external cables, and power ON the engine and peripheral devices.

Related topics:

- "Before you begin replacing components" on page 426
- "Safety information" on page 418

Replacing the hard disk backplane

Steps:

Perform the following steps to remove the hard disk backplane:

1. Read the safety information.
2. Power OFF the Model 1RX and any peripheral devices and disconnect all external cables and power cords.
3. Remove the cover.
4. Remove the faulty hard disk backplane:
   a. Pull out the hot-swap hard disk drives just far enough to disengage the hard disk backplane assembly.
   b. Disconnect the two cables from the diagnostics or operator panel card.
   c. Disconnect the hard disk cable from the backplane.
   d. Disconnect the two cables from the back of the CD-ROM drive.

e. Press on the two backplane release tabs and pull out the hard disk backplane.

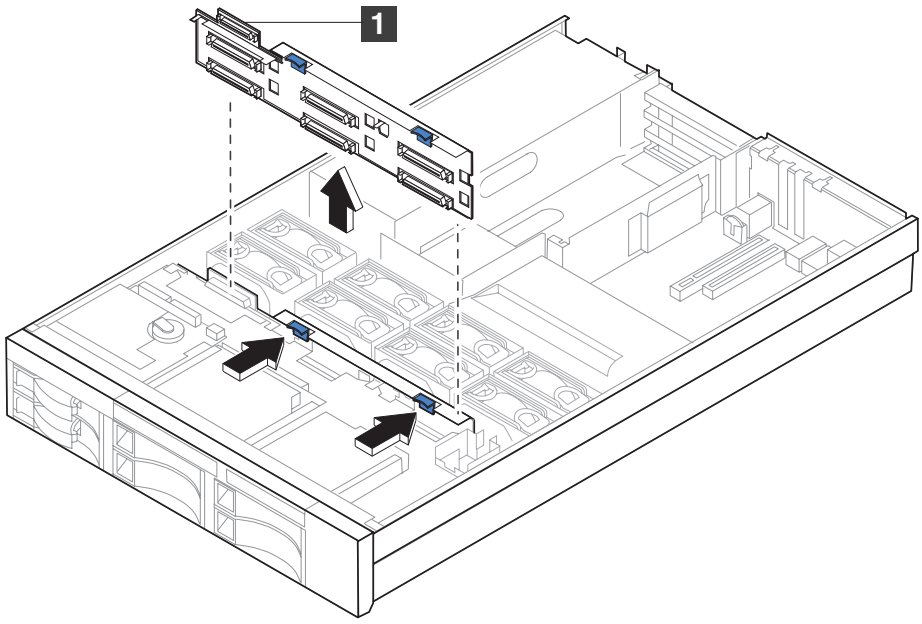

1 Hard disk backplane

f. Disconnect the power cable from the backplane.

5. Repeat step 4 on page 120 in reverse order to install the replacement hard disk backplane.

   Note: When installing the replacement hard disk backplane, align the bayonet on the hard disk cage with the holes in the backplane and engage the left-most tab first.

6. Connect all external cables and power cords, and power ON the Model 1RX and any peripheral devices.

Related topics:

- "Before you begin replacing components" on page 426
- "Replacing the cover and bezel" on page 144
- "Removing the cover and bezel" on page 110
- "Replacing a hot-swap hard disk drive" on page 122
- "Safety information" on page 418

Replacing a diskette drive

Note: The electromagnetic interference (EMI) integrity of the appliance and cooling for the engine are protected by having all drive bays covered or occupied.

Steps:

Perform the following steps to replace the diskette drive:

1. Review the safety information and information about handling static-sensitive devices.
2. Power OFF the engine and peripheral devices. Disconnect all power cords and external cables.

3. Remove the cover.
4. Remove the defective diskette drive:
   a. Unplug the power and data cable attached to the rear of the drive.
   b. Remove the wire spring from the right side of the drive.
   c. Slide the drive out of the bay.
5. Install the new diskette drive:
   a. Align the rails on the drive with the guide rails in the drive bay and push the drive into the bay.
   b. Attach the wire spring.
   c. Plug in the power and data cable at the rear of the drive. The connectors are keyed and can be inserted only one way.
6. Replace the cover, reconnect all power cords and external cables, and power ON the engine and peripheral devices.

Related topics:

- "Before you begin replacing components" on page 426
- "Safety information" on page 418

Replacing a hot-swap hard disk drive

The Model 1RX comes with two 1-inch (26 mm) slim 3.5-inch low voltage differential (LVD) hard disk drives.

Each hot-swap drive has two indicator LEDs. If the amber hard disk drive status LED is lit continuously, that drive is faulty and requires replacement. When the hard disk status LED indicates a drive fault, you can replace a hot-swap drive without powering off the engine.

The following illustration shows how to install a hot-swap hard disk drive.

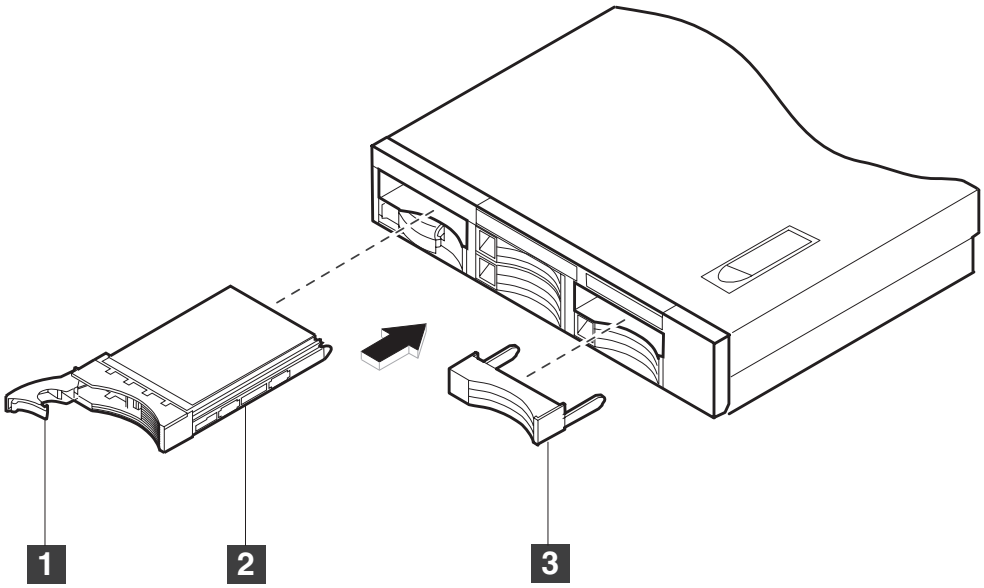

1 Drive handle
2 Drive-tray assembly
3 Filler panel

Notes:

1. All hot-swap drives being used in the Model 1RX should have the same data transfer rate. Mixing hard disk drives with different data transfer rates will cause all drives to operate at the lower speed.
2. The SCSI ID assigned to each bay is printed on the bezel.

Attention: To maintain proper system cooling, do not operate the server for more than 10 minutes without either a hard disk drive or a filler panel installed in each drive bay.

Steps:

Perform the following steps to replace a hot-swap hard disk drive in a drive bay:

1. Review the safety information.
2. Remove the defective hot-swap hard disk drive:
   a. Pull the drive-tray handle to the open position (perpendicular to the hard drive) to disengage the hard disk drive connector from the backplane at the back of the drive bay.
   b. Wait approximately 30 seconds to allow the hard disk drive to spin down.
   c. Pull the hot-swap hard disk drive assembly from the hot-swap drive bay.
3. Install the hard disk drive in the hot-swap bay:
   a. Ensure that the drive handle is open (that is, perpendicular to the hard disk drive).
   b. Align the rails on the drive-tray assembly with the guide rails in the drive bay.
   c. Gently push the drive-tray assembly into the bay until the hard disk drive connects to the backplane.
   d. Push the tray handle to the closed (locked) position until it locks the hard disk drive in place.
4. Check the hard disk drive status LED to verify that the hard disk drive is operating correctly. The Hard-disk-drive status LED is located to the immediate right of the hard disk drive.
   - When the amber LED is lit continuously, the hard disk drive is faulty.
   - When the green LED flashes rapidly (three flashes per second), the controller is identifying the hard disk drive.

Related topics:

- "Before you begin replacing components" on page 426
- "Safety information" on page 418

Replacing a hot-swap fan

The Model 1RX comes with eight hot-swap-fan assemblies making the fans fully redundant. You do not need to power off the Model 1RX to replace a hot-swap-fan assembly.

Attention: To ensure correct operation, if a fan fails, replace it.

The following illustration shows the location of the hot-swap fans.

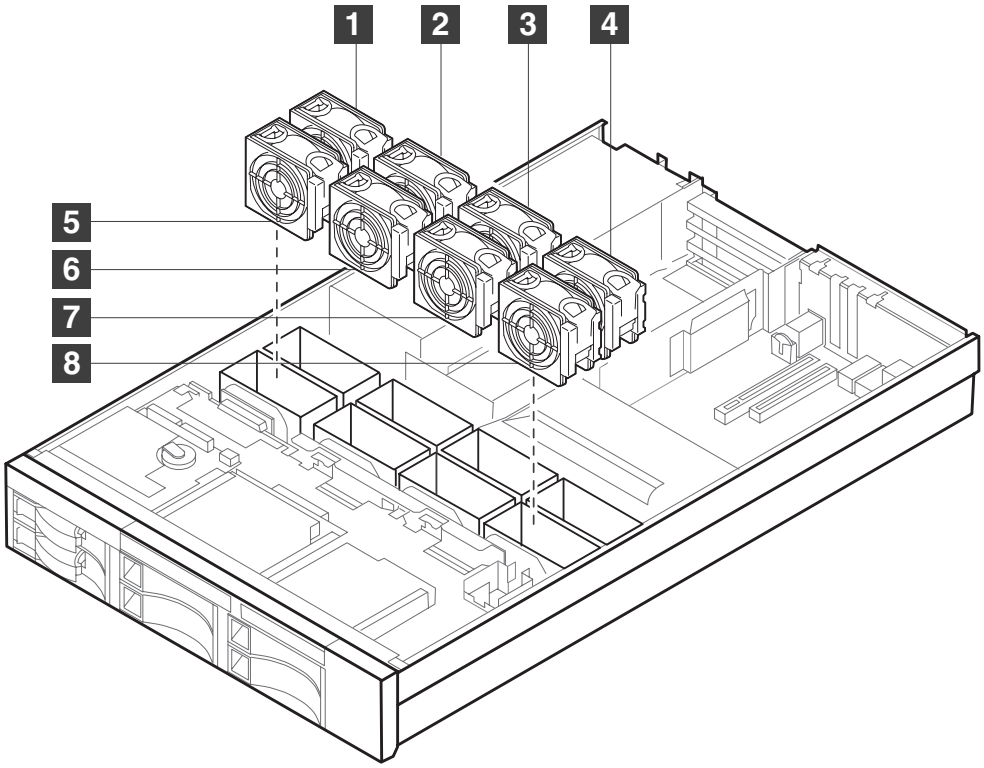

1 Fan 1
2 Fan 2
3 Fan 3
4 Fan 4
5 Fan 5
6 Fan 6
7 Fan 7
8 Fan 8

Steps:

Perform the following steps to replace a hot-swap-fan assembly:

1. Read the safety information.
2. Remove the cover.

   Attention: To ensure proper system cooling, do not remove the top cover for more than 30 minutes during this procedure.
3. Determine which fan assembly to replace by checking the LED on each fan. The LED on the failing fan assembly will be lit.
4. Place your fingers into the two handles on the top of the faulty fan assembly.
5. Squeeze the handles together and lift the fan out of the Model 1RX.
6. Orient the replacement fan to ensure that the LED on top of the fan is to the right of the Model 1RX (facing the front of the box).
7. Push the replacement fan assembly into the Model 1RX until it clicks into place.
8. Replace the cover.

Related topics:

- "Before you begin replacing components" on page 426

- "Replacing the cover and bezel" on page 144
- "Removing the cover and bezel" on page 110
- "Safety information" on page 418

Replacing a hot-swap power supply

The Model 1RX has two hot-swap power supplies for redundancy.

Attention: During normal operation, each power-supply bay must have either a power supply or power-supply blank installed for proper cooling.

DANGER

Do not try to open the covers of the power supply module. (D02)

DANGER

Do not plug the power cable into the power supply module until the enclosure is completely installed, its retaining screws are tightened, and all signal cables are connected. (D03)

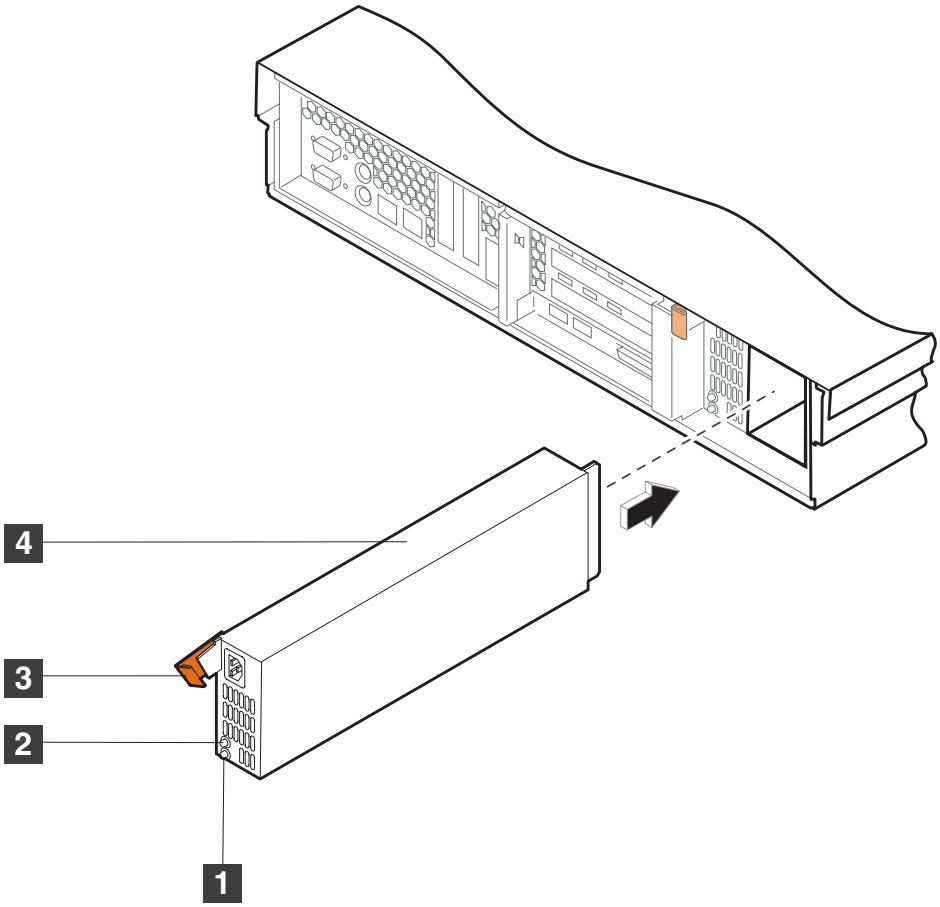

1 DC power LED (green)
2 AC power LED (green)
3 Handle (open position)
4 Power supply

Notes:

1. The Model 1RX provides power redundancy and the ability to hot-swap certain components; therefore, you do not need to power off the engine to install hot-swap power supplies. However, if the load on your system requires the capacity of all installed power supplies, you do not have redundancy or the ability to hot-swap power supplies, and must power off the system before removing any of the power supplies.
2. When replacing a power supply, do not remove the power supply until you are ready to install its replacement.

Attention: If you remove a power supply while the engine is running and you do not have power redundancy, your appliance will abruptly cease to function.

Steps:

Perform the following steps to replace a hot-swap power supply:

1. Read the safety information.
2. Remove the cover.
3. Determine which power supply has failed by checking the Power Supply LEDs on the Light Path Diagnostics panel on the system planar.
   - If the Non Redundant LED is on, you do not have redundancy. Power OFF the engine and peripheral devices.
   - If the Non Redundant LED is not on, you have redundancy and do not need to power off the engine.
4. Remove the faulty power supply:
   a. Disconnect the power cable from the back of the failing power supply. Do not disconnect the power cables from any other power supplies.
   b. Press firmly on the tab behind the handle on the failing power supply and then pull out on the handle to release the power supply.
   c. Gently slide out the power supply.
5. Install the replacement power supply in the bay:
   a. Place the handle on the power supply in the open position (that is, perpendicular to the power supply) by pinching the side clip and pressing down on the handle.
   b. Slide the power supply into the chassis and gently push the power supply into the bay until it connects to the power supply backplane connector.
   c. Press on the handle to seat the power supply in the bay.
6. Plug the power cord for the new power supply into the power-cord connector on the rear of the new power supply.

   The following illustration shows the rear of the server.

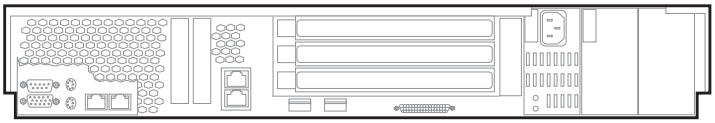

7. If you powered off the engine and any peripheral devices, power ON these devices.
8. Verify that the dc power LED and AC power LED on the power supply are lit, indicating that the power supply is operating correctly.

Related topics:

- "Before you begin replacing components" on page 426
- "Safety information" on page 418
- "Model 1RX rear view" on page 366

Replacing a memory module

The Model 1RX comes with four 1-GB dual inline memory modules (DIMMs) installed on the system board in the DIMM connectors, for a total of 4-GB of memory. This is the maximum amount of memory that the Model 1RX supports.

Memory modules are installed in matched pairs in the order shown in the following table. Both DIMMs in a pair must be the same size, speed, type, and technology, although the second pair of DIMMs does not have to be the same size, speed, type, and technology as the first pair. You can mix compatible DIMMs from various manufacturers.

Table 6. DIMM order of installation

Pair	DIMM connectors
First	1 and 2 (J1, J2)
Second	3 and 4 (J3, J4)

Notes:

1. The amount of usable memory will be reduced depending on the system configuration. The BIOS will display the total amount of installed memory and the amount of configured memory.
2. Install only 100 MHz, 2.5 V, 184-pin, double-data-rate (DDR), PC2100, registered synchronous dynamic random-access memory (SDRAM) with error correcting code (ECC) DIMMs. These DIMMs must be compatible with the latest PC2100 (and downward compatible with PC1600) SDRAM Registered DIMM specification, which is available from www.jedec.org on the World Wide Web.
3. Save new configuration information only if you replace a faulty DIMM that was marked as **Disabled** in the Configuration/Setup Utility program **Memory Settings** menu. In this case, re-enable the memory row in the Configuration/Setup Utility program or reload the default memory settings.

Attention: When you handle devices sensitive to electrostatic discharge (ESDs), take precautions to avoid damage from static electricity.

DANGER

Electrical current from power, telephone and communication cables is hazardous. (D05)

To avoid a shock hazard:

- **Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.**
- **Connect all power cords to a properly wired and grounded electrical outlet.**
- **Connect to properly wired outlets any equipment that will be attached to this product.**
- **When possible, use one hand only to connect or disconnect signal cables.**
- **Never turn on any equipment when there is evidence of fire, water, or structural damage.**
- **Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.**
- **Connect and disconnect cables as described in the following table when installing, moving, or opening covers on this product or attached devices.**

To Connect	To Disconnect
1. Turn everything OFF. 2. First, attach all cables to devices. 3. Attach signal cables to connectors. 4. Attach power cords to outlet. 5. Turn device ON.	1. Turn everything OFF. 2. First, remove power cords from outlet. 3. Remove signal cables from connectors. 4. Remove all cables from devices.

CAUTION:
The power control button on the device and the power switch on the power supply do not turn off the electrical current supplied to the device. The device also might have more than one power cord. To remove all electrical current from the device, ensure that all power cords are disconnected from the power source. (C26)

Note: For translations of this safety notice, refer to *Translated Safety Notices*, GA67-0043.

The following illustration shows how to replace DIMMs on the system board.

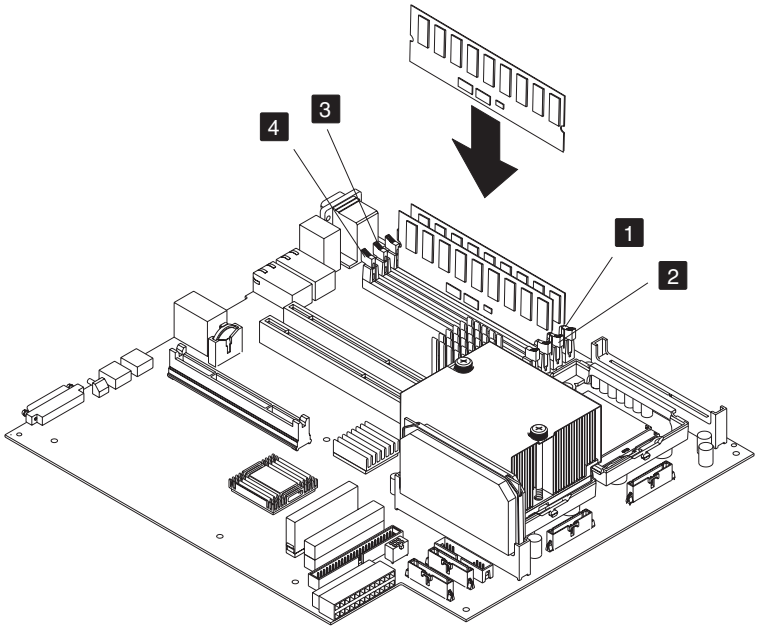

1 DIMM 1 (J1)
2 DIMM 2 (J2)
3 DIMM 3 (J3)
4 DIMM 4 (J4)

Steps:

Perform the following steps to replace a memory module (DIMM):

1. Review the safety information and the documentation that comes with the DIMM.
2. Power OFF the engine and any peripheral devices, and disconnect all external cables and power cords.
3. Remove the cover.
4. Remove the defective DIMM by pressing firmly on the white retaining clips on either side of the DIMM. The DIMM will pop out of the connector.
5. Touch the static-protective package that contains the DIMM option to any unpainted metal surface on the engine, and remove the DIMM from the package. This discharges any static electricity from the package and from your body.
6. Ensure that the DIMM retaining clips are in the open position.

   Attention: To avoid breaking the retaining clips or damaging the DIMM connectors, handle the clips gently.
7. Turn the DIMM to ensure that the pins align correctly with the connector. Notice the notches in the DIMM edge that key the DIMM to the connector.
8. Insert the DIMM into the connector by pressing on one end of the DIMM and then on the other end of the DIMM. Be sure to press straight into the connector. Be sure that the retaining clips snap into the closed position.

   Note: If a gap exists between the DIMM and the retaining clips, the DIMM has not been installed correctly. In this case, open the retaining clips and remove the DIMM; then, reinsert the DIMM.
9. Repeat these steps for the remaining DIMM.

10. Replace the cover, connect external cables and power cords, and power ON the engine and any peripheral devices.

Related topics:

- "Before you begin replacing components" on page 426
- "Handling static-sensitive devices" on page 424
- "Replacing the cover and bezel" on page 144
- "Removing the cover and bezel" on page 110
- "Safety information" on page 418
- "Starting the Configuration/Setup Utility" on page 106

Replacing a microprocessor

The Model 1RX comes with two microprocessors, which enables each engine to operate as a symmetric multiprocessing (SMP) server. With SMP, SAN File System distributes the processing load between the microprocessors.

Attention:

- A startup (boot) microprocessor must always be installed in socket J7 on the system planar.
- To ensure proper operation, the microprocessors must have the same cache size and type, and the same clock speed. Microprocessor internal and external clock frequencies must be identical.
- When you handle devices that are sensitive to electrostatic discharge, take precautions to avoid damage from static electricity.

Notes:

1. Before you install a new microprocessor, refer to the support Web site, at http://www.ibm.com/storage/support, to determine whether a BIOS upgrade is required.
2. The microprocessor speeds are automatically set for this engine; therefore, you do not need to set any microprocessor frequency-selection jumpers or switches.

The following illustration is a simplified layout of the microprocessor connector locations and other microprocessor-related components on the system planar.

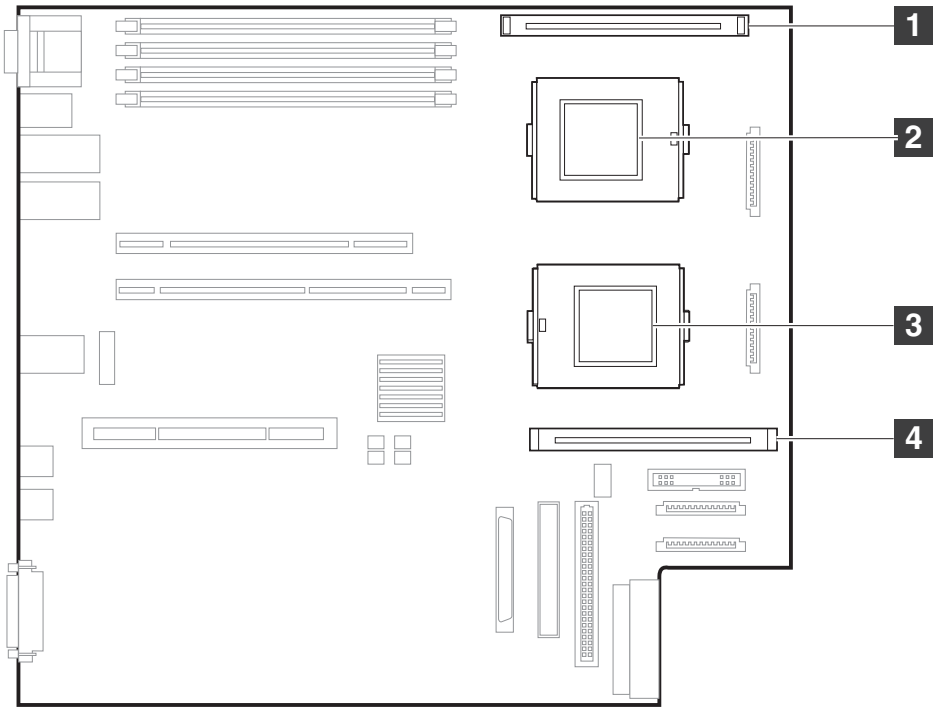

1 VRM 2 connector (J5)
2 Microprocessor 2 (J6)
3 Microprocessor 1 (J7)
4 VRM 1 connector (J8)

The following illustration shows how to replace a microprocessor on the system planar.

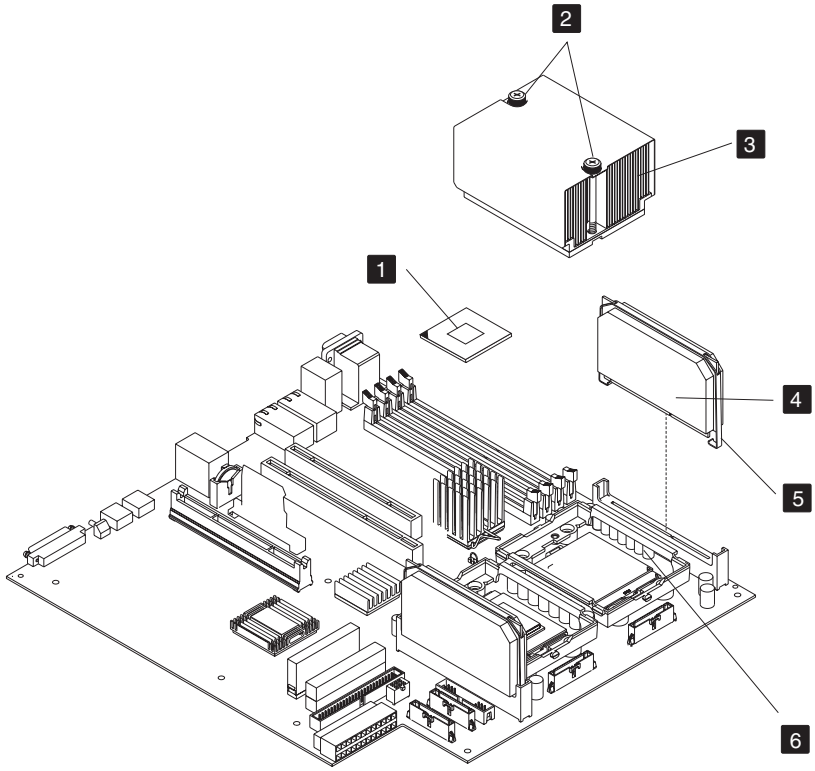

1 Microprocessor
2 Captive screws
3 Heat sink
4 VRM
5 VRM clip
6 Microprocessor release lever

Steps:

Perform the following steps to replace a microprocessor:

1. Review the safety information.
2. Power OFF the engine and peripheral devices, and disconnect all external cables and power cords.
3. Remove the cover.

4. Remove the air baffle covering the microprocessor section of the system planar by pressing on the tab and lifting it up. Set the air baffle aside.

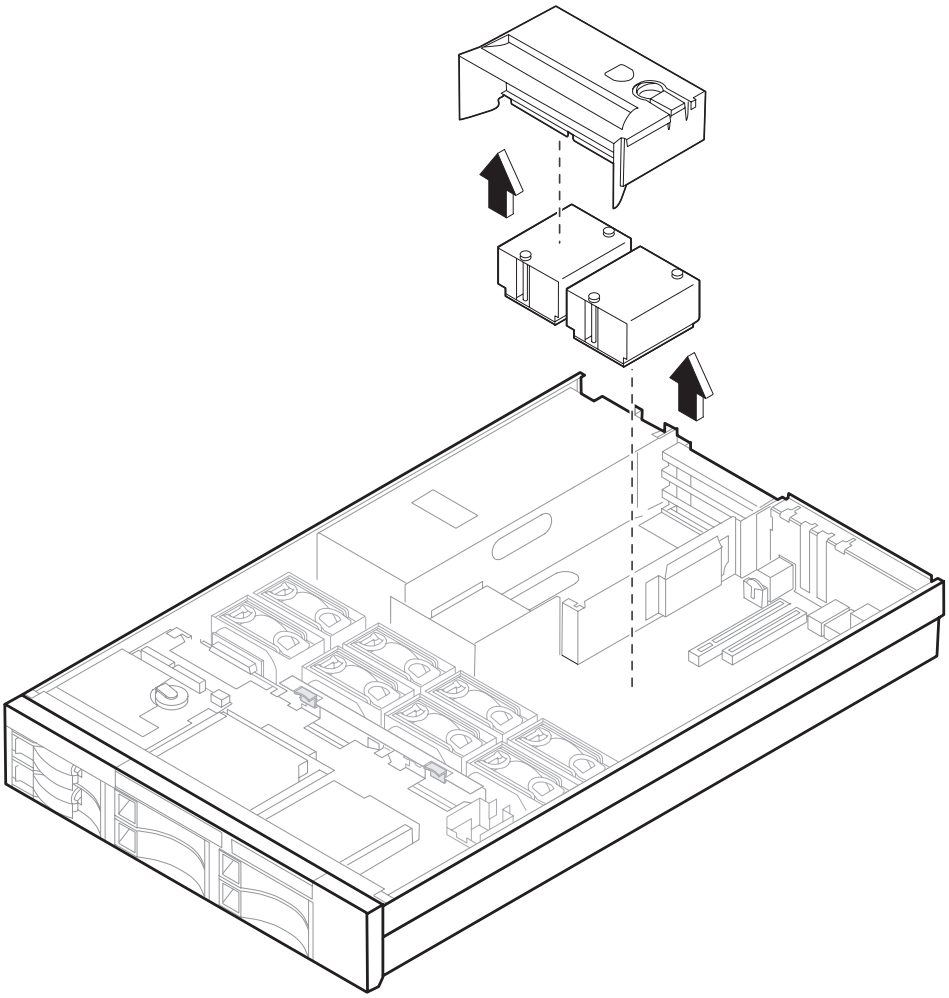

5. Remove the faulty VRM:
   a. Remove the retention spring by pressing down on each corner and lifting the spring away from the VRM.
   b. Lift the VRM out of the connector.
6. Remove the heat sink by loosening the screws and lifting the heat sink.
7. Remove the faulty microprocessor by raising the release lever and lifting up the microprocessor.
8. Install the replacement microprocessor:
   a. Touch the static-protective package containing the new microprocessor to any *unpainted* metal surface on the engine; then, remove the microprocessor from the package.
   b. Remove the protective tape from the microprocessor socket, if the tape is present.
   c. Center the microprocessor over the microprocessor socket. Align the triangle on the corner of the microprocessor with the triangle on the corner of the socket and carefully set the microprocessor in the socket.

Attention:

- Do not use excessive force when pressing the microprocessor into the socket.
- Make sure that the microprocessor is oriented and aligned correctly with pin number 1 in the socket before you try to close the lever. The following illustration shows the alignment marks.

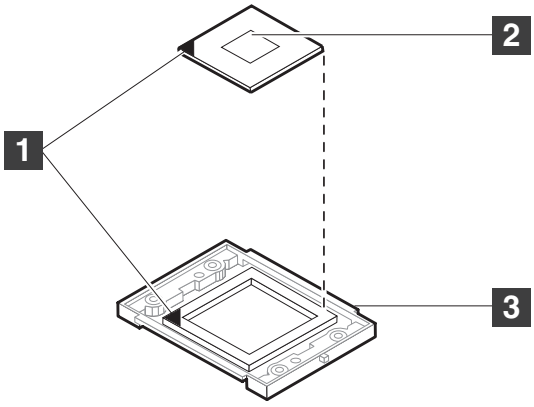

1 Alignment mark
2 Microprocessor
3 Microprocessor socket

d. Carefully close the lever to secure the microprocessor in the socket.

9. Install a heat sink on the microprocessor:

Attention:

- Do not set the heat sink down after you remove the plastic cover.
- Do not touch the thermal grease on the bottom of the heat sink. Touching the thermal grease will contaminate it. If the thermal grease on the microprocessor or heat sink becomes contaminated, contact your service technician.

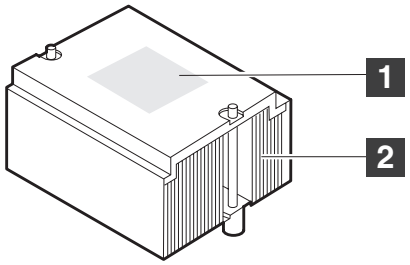

1 Thermal grease
2 Heat sink

a. Remove the protective plastic cover from the bottom of the heat sink.

b. Align and place the heat sink on top of the microprocessor with the thermal grease side down. Press firmly on the heat sink.

c. Secure the heat sink to the microprocessor by tightening the captive screws. Alternate between the two captive screws in the heat sink while tightening them.

Attention: If you do not alternate the tightening of the two captive screws, you might damage the microprocessor. Ensure that the screws are completely tightened; otherwise, the engine will not function properly.

10. Install the replacement VRM (see the illustration at the beginning of these instructions for the location of the connectors):
    a. Center the VRM over the connector. Make sure that the VRM is oriented and aligned correctly.
    b. Carefully but firmly push down the VRM to seat it in the connector.
    c. Insert one end of the retainer clip on the side of the connector and fit the retainer clip over the VRM. Insert the free end of the retainer clip on the other side of the connector.

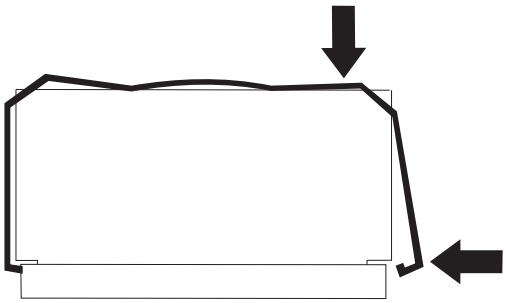

11. Install the air baffle over the microprocessor section of the system planar by aligning the notch on the side of the shuttle with the tab on the air baffle.
12. Replace the cover, connect external cables and power cords, and power ON the engine and peripheral devices.

Related topics:

- "Before you begin replacing components" on page 426
- "Handling static-sensitive devices" on page 424
- "Replacing the cover and bezel" on page 144
- "Removing the cover and bezel" on page 110
- "Safety information" on page 418

Replacing the diagnostics/operator panel card

Steps:

Perform the following steps to replace the diagnostics/operator panel card:

1. Read the safety information.
2. Power OFF the Model 1RX and any peripheral devices, and disconnect external cables and power cables.
3. Remove the cover.
4. Remove the diagnostics/operator panel card:
    a. Disconnect the three cables from the card.
    b. Remove the two screws from the card.

c. Pull out the diagnostics/operator panel card.

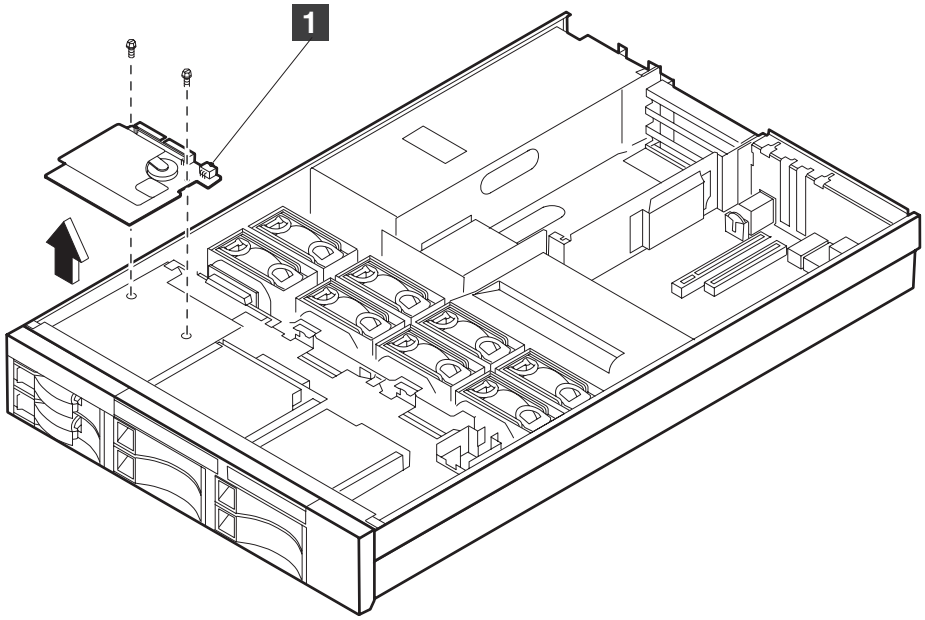

1 Diagnostics/operator panel card

5. Repeat step 4 on page 136 in reverse order to replace the diagnostics/operator panel card.
6. Replace the cover, connect external cables and power cables, and power ON the Model 1RX and any peripheral devices.

Related topics:

- "Before you begin replacing components" on page 426
- "Replacing the cover and bezel" on page 144
- "Removing the cover and bezel" on page 110
- "Safety information" on page 418

Replacing the power-supply cage

Attention: To ensure correct operation, do not remove the air-baffle assembly from the Model 1RX except when installing or removing the components that are located under the air baffle.

Steps:

Perform the following steps to replace the power-supply cage:

1. Read the safety information.
2. Power OFF the Model 1RX and any peripheral devices, and disconnect external cables and power cables.
3. Remove the cover.
4. Remove the faulty power-supply cage:
   a. Remove the power supplies.
   b. Remove the air baffle by pressing the integrated latches and lifting it up.
   c. Partially remove the DASD backplane.
   d. Remove the PCI riser card assembly.

e. Disconnect the four power-supply cage cables.

f. Remove the three screws from the power-supply cage by lifting up the side where the cables connect and pulling it out.

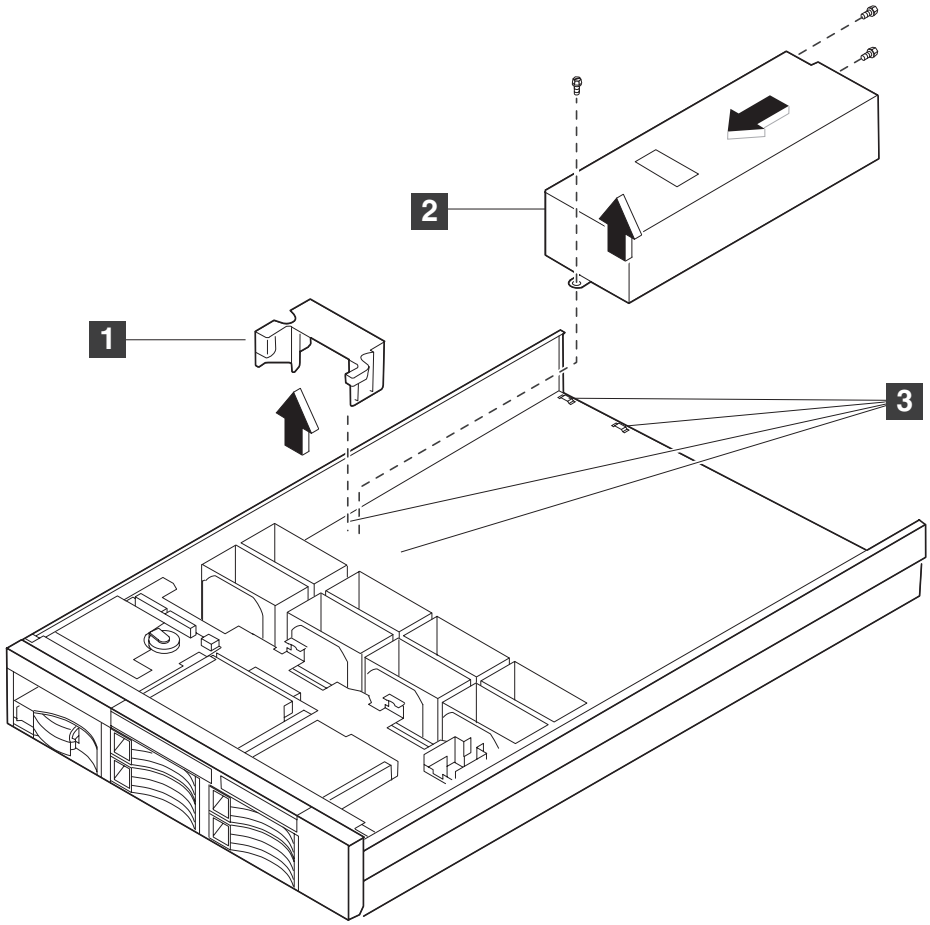

1 Power-supply air baffle
2 Power-supply cage
3 Raised tabs

g. Pull out the power-supply cage.

5. Repeat step 4 on page 137 in reverse order to install the replacement power-supply cage.

   Note: Tilt the power-supply cage during installation to ensure that you insert the rear of the cage first.

6. Replace the cover, connect external cables and power cables, and power ON the Model 1RX and any peripheral devices.

Related topics:

- "Before you begin replacing components" on page 426
- "Handling static-sensitive devices" on page 424
- "Replacing the cover and bezel" on page 144
- "Removing the cover and bezel" on page 110
- "Adding or replacing an adapter" on page 111
- "Replacing the hard disk backplane" on page 120
- "Replacing a hot-swap power supply" on page 125

- "Safety information" on page 418

Replacing an RSA II adapter

Prerequisites:

A Phillips screwdriver may be required.

Steps:

The following illustration shows how to install an RSA II adapter.

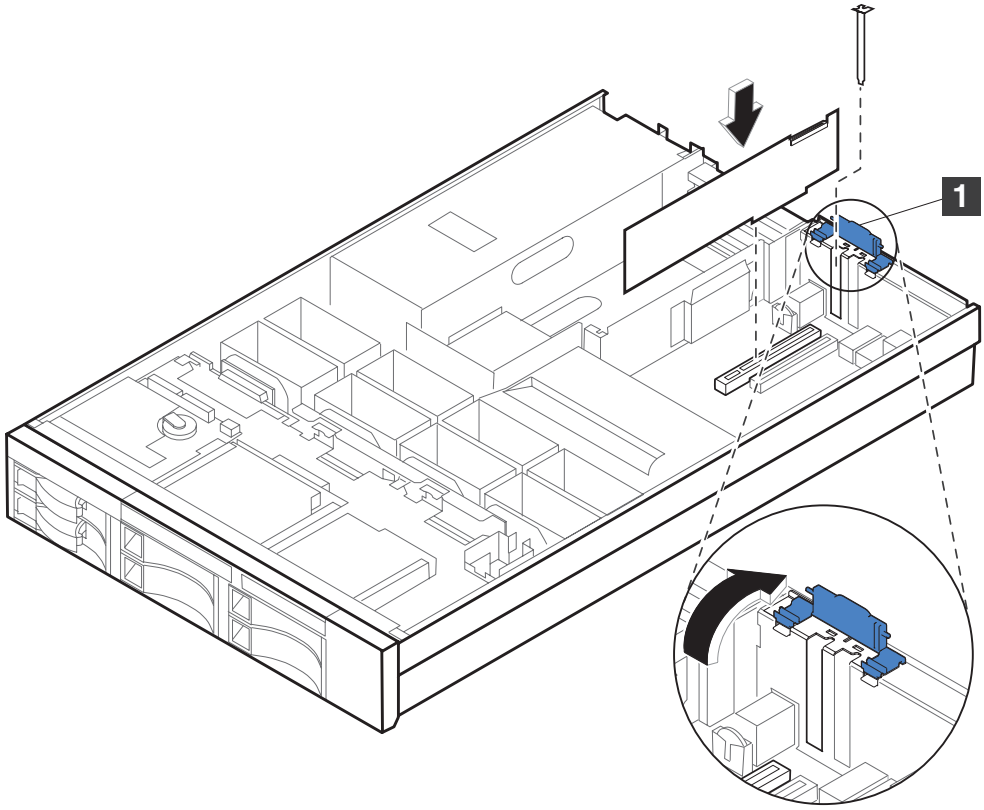

1 Adapter retainer

Perform the following steps to replace an RSA II adapter:

1. Review the safety information and information about handling static-sensitive devices.
2. Verify that the Metadata server is offline.
3. Power OFF the engine and any peripheral devices, and disconnect all external cables and power codes.
4. Remove the engine from the rack.
5. Remove the cover.

6. Remove the PCI riser-card assembly:

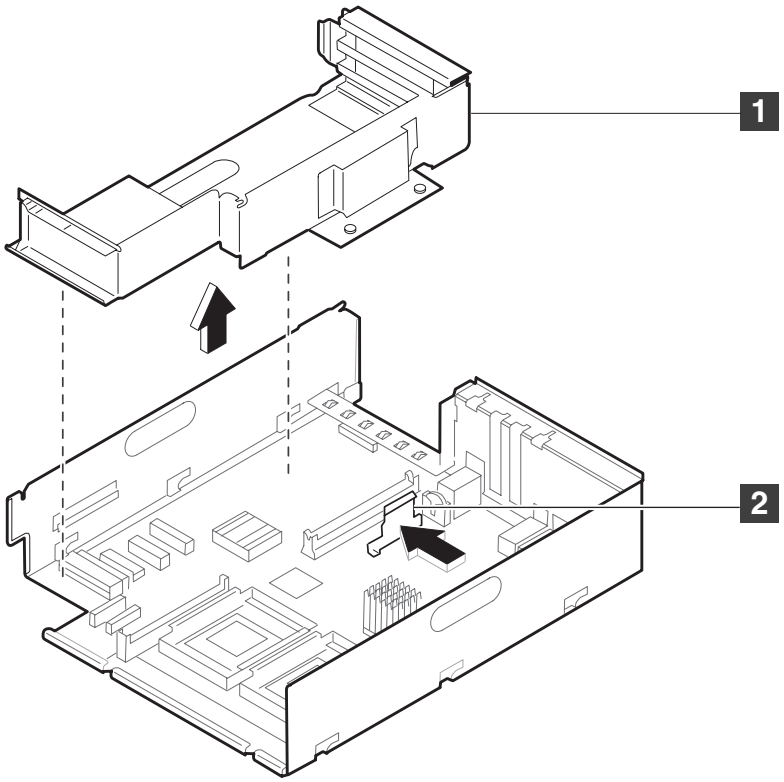

1 PCI-riser card assembly
2 Retention latch

   a. If the PCI riser-card assembly has a latch, release the latch. Otherwise, loosen the two captive screws to release the PCI riser-card assembly.
   b. Lift the assembly from the engine.
   c. Place the assembly on a flat, level surface.
7. Remove the defective adapter.
   a. Lift up the blue adapter-retainer.
   b. Carefully lift out the adapter.
   c. Disconnect the adapter cable by squeezing the retention latches and pulling away from the adapter.
8. Install the new RSA II adapter.
   a. Touch the static-protective package containing the adapter to an unpainted metal surface, such as the expansion-slot cover on the engine. This discharges any static electricity from the package and from your body.

      Attention: Avoid touching the components and gold-edge connectors on the adapter.
   b. Remove the new RSA II adapter from the static-protective package.
   c. Carefully grasp the adapter by its top edge or upper corners and align it with the PCI slot on the system board.

      Attention: When you install the adapter in the engine, be sure that it is completely and correctly seated in the system-board connector before you apply power. Incomplete insertion might cause damage to the system board or the adapter.
   d. Press the adapter firmly into the PCI slot.
9. Connect the 20-pin ribbon cable.

a. Connect one end of the 20-pin ribbon cable to the system-management connector on the RSA II adapter. The 20-pin ribbon cable connector ends are keyed for proper installation. Make sure the contrasting color stripe on the cable is closest to the adapter edge that has the gold edge connectors.

   Attention: Route cables so that they do not block the flow of air from the fans.

b. Connect the other end of the 20-pin ribbon cable to the RSA connector (sometimes called the system-management connector) on the system board or riser card:

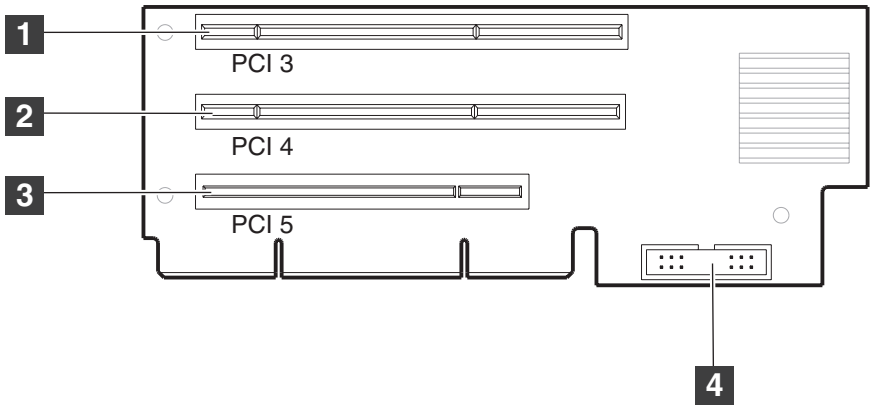

1 PCI-X slot 3 64-bit 3.3V 133 MHz (PCI 3)
2 PCI-X slot 4 64-bit 3.3V 133 MHz (PCI 4)
3 PCI-X slot 5 32–bit 5V 33 MHz (PCI 5)
4 System management connector

10. Replace the cover on the engine.
11. Connect the mini-USB cable to the mini-USB connector on the RSA II card and to the USB connector on the engine.
12. Install the engine back in the rack.
13. Connect the Ethernet cable to the Ethernet connector on the RSA II adapter.
    a. Connect one end of a Category 3 or Category 5 Ethernet cable to the Ethernet connector on the RSA II adapter. Connect the other end of the Ethernet cable to your network.
    b. Check the RSA II Ethernet LEDs and power and the adapter activity LEDs.
       1 Ethernet Link LED. The green Ethernet link LED is lit when the Ethernet cable is connected properly.
       2 Ethernet activity LED. The green Ethernet activity LED is flashing when there is Ethernet activity.
    c. Connect the video cable of your video device to the video connector on the RSA II adapter. This is required for video.
14. Reconnect the external cables and power cords that you disconnected.
15. Power ON the engine.
16. Connect the ASM breakout cable by connecting a Category 3 or Category 5 cable to one of the ASM RS-485 connectors on the ASM breakout cable.
17. Reconfigure the RSA II IP address:
    a. Open a browser from the remote console on the master console and point it to the default URL for the RSA adapter

       `http://192.168.72.125/`

    b. Log on to the RSA II interface using a valid RSA user name and password.
    c. Click **Server→ASM Control→Network Interface**.

d. From the Ethernet page, enter the correct IP address of the RSA II card (based on the engine in which it is installed).
e. Click **Save** to save the configuration.

Post-processing requirements:

After you have completed the installation of the RSA II card, you may need to update one or more of the following software and firmware:

- Video adapter BIOS code
- Engine BIOS code
- Engine diagnostics code
- RSA II firmware code
- RSA II device drivers

These software and firmware files are contained in one system service package that is available from the TotalStorage 4146 support Web site at: www.ibm.com/storage/support/4146.

Related topics:

- "Accessing the RSA II adapter" on page 24

Replacing the system planar and shuttle

Note: When replacing the system planar, you must either update the Model 1RX with the latest firmware or restore the preexisting firmware that the customer provides on a diskette or CD image.

Steps:

Perform the following steps to replace the system planar and shuttle assembly:

1. Read the safety information.
2. Power OFF the Model 1RX and any peripheral devices, and disconnect all external cables and power cords.
3. Remove the cover.
4. Remove the faulty system planar and shuttle assembly:
   a. Disconnect and remove adapters.
   b. Disconnect all cables from the system planar.
   c. Remove all microprocessors and VRMs and set them aside on a static-protected surface for reinstallation.
   d. Remove the memory modules and set them aside on a static-protected surface for reinstallation.

e. Release the shuttle-locking latch and remove the system planar and shuttle assembly.

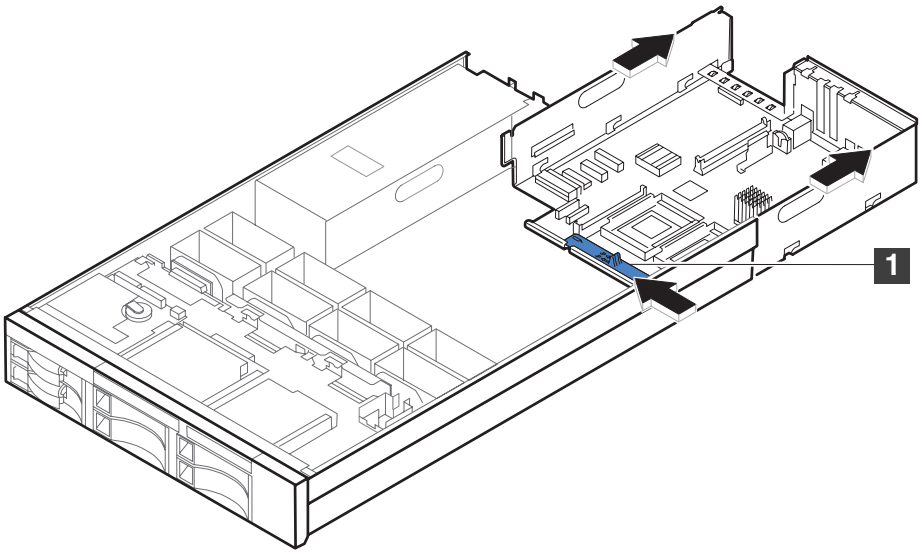

1 Shuttle locking latch

5. Install the replacement system planar and shuttle assembly:
   a. Slide the system planar and shuttle assembly into the Model 1RX and close the shuttle-locking latch.
   b. Remove the air baffle covering the microprocessors.
   c. Install the memory modules.
   d. Install the microprocessors and VRMs.
   e. Connect all cables from the system planar.
   f. Install and connect the adapters.
6. Replace the cover, connect all external cables and power cords, and power ON the Model 1RX and any peripheral devices.

Related topics:

- "Before you begin replacing components" on page 426
- "Replacing the cover and bezel" on page 144
- "Removing the cover and bezel" on page 110
- "Adding or replacing an adapter" on page 111
- "Replacing a memory module" on page 127
- "Replacing a microprocessor" on page 131
- "Safety information" on page 418

Replacing the cover and bezel

Attention: For proper cooling and airflow, replace the cover before or soon after powering on the engine. Operating the engine for extended periods of time (over 30 minutes) with the cover removed might damage components.

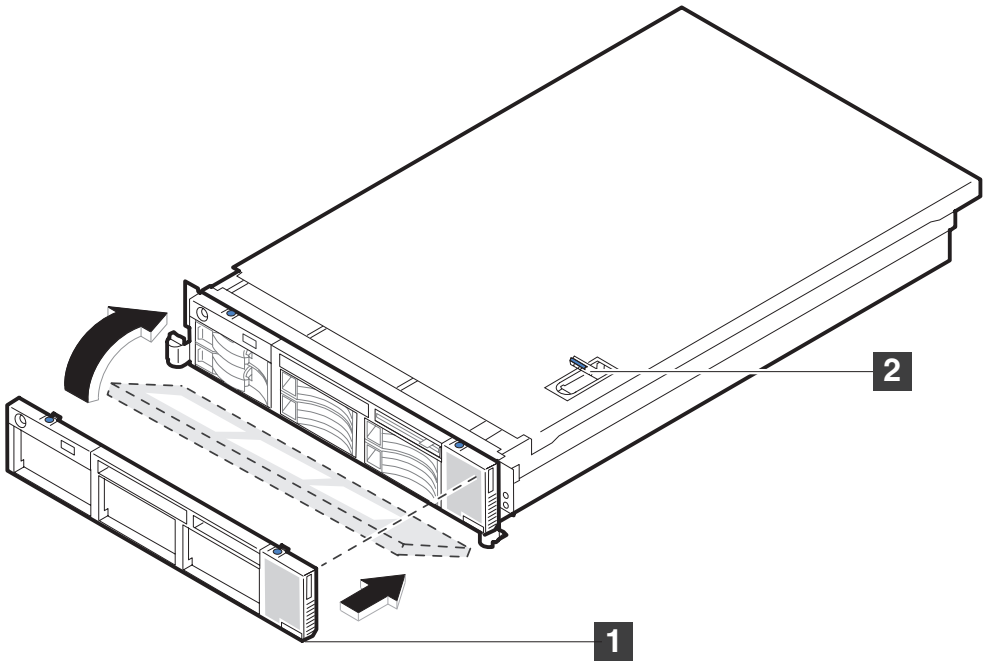

1 Cover-release latch
2 Bezel

Steps:

Perform the following steps to replace the top cover:

1. Review the safety information.
2. Place the cover-release latch in the open (up) position.
3. Insert the bottom tabs of the left-side cover into the matching slots in the chassis.
4. Press down on the cover-release latch to lock the cover in place.

Perform the following steps to replace the bezel:

1. Align the hooks on the bottom of the trim bezel with the Model 1RX.
2. Press the top of the bezel toward the Model 1RX until it clicks into place.

Related topics:

- "Safety information" on page 418
- "Removing the cover and bezel" on page 110

Chapter 12. Disaster recovery

The following steps are used to restore functionality of the SAN File System following an unexpected outage:

Note: This section assumes that you followed the steps for managing disaster recovery that are described in the *Planning, Installation and Configuration Guide*.

1. Restore the hardware and operating system for the master console.
2. Restore the hardware and operating system for each of the engines in the cluster.
3. Restore SAN connectivity (LUN configuration).
4. Restore the SAN File System software (both for the Metadata server and the Administrative server).
5. Restore the SAN File System cluster configuration.
6. Restore the SAN File System metadata.
7. Restore the SAN File System clients.
8. Restore the SAN File System user files.

Related topics:

- Backup and restore overview in the *Planning, Installation and Configuration Guide*
- "Restoring the master console"
- "Restoring the engine hardware and operating system" on page 146
- "Restoring SAN connectivity" on page 146
- "Restoring SAN File System software" on page 147
- "Restoring SAN File System cluster configuration" on page 147
- "Restoring SAN File System metadata" on page 149
- "Restoring SAN File System clients" on page 150
- "Restoring SAN File System user data" on page 150

Restoring the master console

Steps:

Perform these steps to restore the hardware and operating system for the master console.

1. Determine if the hardware for the master console is working properly. If so, review information about recovering the hard drives (if necessary) as well as recovering the software.
2. If the hardware for the master console is not working properly,
   - Refer to the refer to the *IBM @server xSeries 305 Hardware Maintenance Manual and Troubleshooting Guide* to resolve problems with the hardware.
   - Refer to the *Planning, Installation and Configuration Guide* for information about installing the master console.

Related topics:

- "Resolution procedures" on page 83

- Chapter 9, “Troubleshooting the master console”, on page 83

Restoring the engine hardware and operating system

Steps:

Perform these steps to restore the hardware and operating system for each engine in the SAN File System cluster.

1. Verify that there is no damage to the hardware and that the engine boots properly. If you suspect a problem with any of the hardware components, troubleshoot an engine to resolve the problem.
2. Verify that there is no damage to the master console and that it boots properly.
   - If you suspect a problem with any of the hardware components in the master console, refer to the to the *IBM @server xSeries 305 Hardware Maintenance Manual and Troubleshooting Guide* to resolve those problems.
   - If you suspect a problem with the software or the hard disk drive, troubleshoot the master console to resolve the problem.
3. From the master console, point the Web browser to the URL of the RSA II adapter on the engine and access the RSA II adapter to set up a remote console to the engine. This interface allows you to use the master console as your display and keyboard for the engine.

   Note: Instead of using the RSA II Web interface from the master console, you can directly attach a keyboard and display to the engine. However, make sure that you attach the display to the VGA port of the RSA II card on the engine and not to the video port on the engine itself.
4. Determine if the boot drives for each engine hosting a Metadata server still have an intact SAN File System configuration and executable files (undamaged).
5. If there are corrupt or damaged configuration and executable files, attempt to recover the damaged files from the mirrored boot drive. If you cannot recover the damaged files from the mirrored boot drive:
   a. Load the Disaster Recovery CD into the CD-ROM drive on the engine.
   b. Reboot the engine using one of the following methods:
      - Open a bash shell prompt and enter **init 6**.
      - Press the Reset button on the front panel of the engine.
      - Power off the engine and then power it back on.
   c. When you receive a warning prompt that the entire hard drive will be overwritten, respond by entering **y**.
   d. After the operating system has been reloaded, the engine will eject the Disaster Recovery CD and automatically reboot.

Related topics:

- “Accessing the RSA II adapter” on page 24
- Chapter 10, “Troubleshooting an engine”, on page 87
- Chapter 9, “Troubleshooting the master console”, on page 83

Restoring SAN connectivity

Steps:

If the system was backed up using the LUN method, perform these steps on each engine in the SAN File System cluster to restore SAN connectivity:

1. Verify that the engines hosting the Metadata servers are connected to the SAN in the same configuration that existed at the point of the last backup operation (make sure the Metadata servers can see the same LUNs that existed prior to the unexpected outage).
2. If the LUN mapping has changed, use the device management tools for the storage subsystem or management tools for the SAN to re-create the old LUN map. After creating the old LUN map, reboot the Metadata server so that the changes to the LUN map are visible to the Metadata server.
3. If LUN contents were lost or corrupted, use the copy services facility of the storage subsystem to restore all LUN data (both metadata and user file data).

If the system was backed up using the API method, perform these steps on each engine in the SAN File System cluster to restore SAN connectivity.

1. Verify that the engines hosting the Metadata servers are connected to the SAN in the same configuration that existed at the point of the last backup operation (make sure the Metadata servers can see the same LUNs that existed prior to the unexpected outage).
2. If the LUN mapping has changed, use the device management tools for the storage subsystem or management tools for the SAN to re-create the old LUN map. You can also choose to restore data onto a new LUN map. However, if you do so, you will have to manually run some of the steps used to restore metadata.

Restoring SAN File System software

Steps:

If you had to use the Disaster Recovery CD-ROM to recover the operating system for the engine, perform these steps on each engine in the SAN File System cluster to restore the Metadata server and Administrative server software.

1. Reinstall the software for the Metadata server.
   a. Make sure that you are logged into the engine as root.
   b. From a shell prompt on the engine, change to the directory where the Metadata server software package is installed.

      cd /usr/tank/packages
   c. Install the Metadata server software package using the following command:

      bash# rpm -ivh *Metadata_server_package_name*.**rpm**
   d. Install the Qlogic software using the following command:

      bash# rpm -i snia_qlogic_hba*rpm
2. Reinstall the software package for the Administrative server.
   a. Install the Administrative server software package using the following command:

      bash# rpm -ivh *Admin_server_package_name*.**rpm**

Restoring SAN File System cluster configuration

Steps:

If the system was backed up using the LUN method, perform these steps to restore the cluster configuration information:

1. If you have previously saved the configuration files to another location, copy these files onto the boot drive for the engine.
   a. Copy Tank.Bootstrap to /usr/tank/server/config.
   b. Copy Tank.Config to /usr/tank/server/config.

   Note: If you have saved any other administrative configuration files, you can reference them when restoring the SAN File System metadata configuration.
2. If the cluster bootstrap file, Tank.Bootstrap, is corrupted or missing, you can attempt to re-create the contents of that file using information from the metadata LUNs:
   a. Use the **tank lsdisklabel –device** command to find the master volume. If you cannot remember which device is your master volume, this is an iterative process of searching all suspected master volume devices until the command indicates you have found a valid master volume.
   b. Use the **tank extractbootrecord** command to regenerate Tank.Bootstrap from the master volume.
   c. Use the **tank resetcluster** command to reinitialize the master volume for subsequent rebuilding of the cluster configuration.
   d. Use the **addserver** command for all subordinate MDS engines to re-create the cluster definition.

If the system was backed up using the API method, perform these steps to restore the cluster configuration information:

1. If you have previously saved the configuration files to another location, copy these files onto the boot drive for the engine.
   a. Copy Tank.Bootstrap to /usr/tank/server/config.
   b. Copy Tank.Config to /usr/tank/server/config.

   Note: If you have saved any other administrative configuration files, you can reference them when restoring the SAN File System metadata configuration.
2. If you suspect that the metadata LUNs are corrupted, you can perform these steps to re-create the cluster definition:
   a. Delete all Tank.Bootstrap and Tank.Config files from your MDS engines.
   b. Start the **tank** binary on your master MDS with the **install** option rather than *normal* option. This will create new Tank.Bootstrap and Tank.Config files on your MDS master. Be sure to specify the same cluster name that was used prior to the disaster:
   c. Now start the master MDS with **tank normal** command.
   d. Use the **addserver** command to add all subordinate MDS engines. This will create new Tank.Bootstrap and Tank.Config files on the subordinates.

Related topics:

- "addserver" on page 322
- "tank lsdisklabel" on page 358
- "tank extractbootrecord" on page 353
- "tank resetcluster" on page 360

Restoring SAN File System metadata

Steps:

If the system was backed up using the API method, perform these steps to the SAN File System metadata:

Note: If the system was backed up using the LUN method, the metadata was restored when you restored SAN connectivity.

1. Verify that all Metadata servers in the cluster are online and that the cluster is running.

   ```
   tanktool lsserver -state online
   ```

   All of the online Metadata servers in the cluster are displayed.
2. Copy the system-metadata disaster-recovery file (and the scripts) that you had previously backed up to /usr/tank/server/DR on the master Metadata server.
3. Use the TankSysCLI.auto script:
   a. Edit the script TankSysCLI.auto for information about how the script is used and any changes that may need to be made to the script.

      ```
      ######################################################################
      # CLI Commands to create Storage Pools, Containers, Service Classes and
      # Policy Sets.
      # These commands need NO manual intervention.
      ######################################################################
      ```

   b. Run the script TankSysCLI.auto.

      ```
      tanktool -script /usr/tank/server/DR/TankSysCLI.auto
      ```

   c. If any errors occurred while running the script, make sure that you resolve those errors before continuing.
4. Use the TankSysCLI.volume script:
   a. Edit /usr/tank/server/DR/TankSysCLI.volume and modify it to match your current SAN settings. It also contains usage information as well as information about any changes that may need to be made to the script.

      ```
      ######################################################################
      # CLI Commands to add Volumes to Storage pools.
      # These commands need manual intervention.
      # The device names were as they appeared during backup.
      # Please make sure that the device names appearing here actually
      # exist and have correct sizes and if not edit the device names to
      # correct values.
      # The System MASTER volume has to be specified in tank install command
      # and therefore has no corresponding CLI.
      # The other System Volumes can either be specified in tank install
      # command, or, added using the CLI command, which appears inside comments
      # forthis reason.
      ######################################################################
      ```

   b. Run the script TankSysCLI.volume.

      ```
      tanktool -script /usr/tank/server/DR/TankSysCLI.volume
      ```

   c. If any errors occurred while running the script, make sure that you resolve those errors before continuing.
5. Use the TankSysCLI.attachpoint script:
   a. Edit /usr/tank/server/DR/TankSysCLI.attachpoint to verify the settings. It also contains usage information as well as information about any changes that may need to be made to the script.

```
#########################################################################
# CLI Commands to attach containers.
# These commands need manual intervention.
# All the "mkdir" and "attachcontainer" commands should be run in the
# order given.
# The "mkdir" command should be run on a client to recreate the directory
# path before running the following attachcontainer CLI commands.
#########################################################################
```

b. If all filesets are attached only to the root directories of other filesets, run the script TankSysCLI.attachpoint.

```
tanktool -script /usr/tank/server/DR/TankSysCLI.attachpoint
```

Note: If you have any filesets attached to directories, you have to reattach them manually.

c. If any errors occurred while running the script, make sure that you resolve those errors before continuing.

6. Grant privileges to those clients that require root or Administrator access to SAN File System using the **chclusterconfig –privclient** command.

Related topics:

• "Administrative commands" on page 158

Restoring SAN File System clients

SAN File System clients are access points to the SAN File System. Therefore, clients are not backed up. To restore a SAN File System clients, you can perform the normal client installation procedure, which is described in the *Planning, Installation, and Configuration Guide*.

Related topics:

• Chapter 12, "Disaster recovery", on page 145

Restoring SAN File System user data

Steps:

If the system was backed up using the API method, perform these steps to restore user data:

Note: If the system was backed up using the LUN method, user data was restored when you restored SAN connectivity.

1. From a client, mount the SAN File System at its usual mount point. The top of the subdirectory tree (the portion of the subdirectory tree that consists of the fileset names) should be visible from the client.
2. Restore the files onto that mount point. Follow the procedures for the backup and recovery application used by the customer to back up the files.
3. If the customer has followed the guidelines in the *Planning, Installation, and Configuration Guide* for backup and recovery, restore files to the Windows filesets from a Windows client and restore files to the AIX filesets from an AIX client.

Related topics:

• Chapter 12, "Disaster recovery", on page 145

Chapter 13. Getting help, service, and information

If you need help, service, technical assistance, or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you.

IBM maintains pages on the World Wide Web where you can get information about IBM products and services and find the latest technical information.

Table 7 lists some of these pages.

Table 7. IBM Web sites for help, services, and information

www.ibm.com/	Main IBM home page
www.ibm.com/storage/	IBM Storage home page
www.ibm.com/storage/support	IBM Support home page

Services available and telephone numbers listed are subject to change without notice.

Software Maintenance

All distributed software licenses include Software Maintenance (software subscription and technical support) for a period of 12 months from the date of acquisition providing a streamlined way to acquire IBM software and assure technical support coverage for all licenses. Extending coverage for a total of three years from date of acquisition may be elected. While your Software Maintenance is in effect, IBM will provide you assistance for your 1) routine, short duration installation and usage (how-to) questions; and 2) code-related questions. IBM provides assistance via telephone and, if available, electronic access, only to your information systems (IS) technical support personnel during the normal business hours (published prime shift hours) of your IBM support center. (This assistance is not available to your end users.) IBM provides Severity 1 assistance 24 hours a day, every day of the year.

Hardware Warranty

For a period of one year, if required, IBM provides repair or exchange service depending on the type of warranty service specified for your machine. An IBM technician will attempt to resolve your problem over the telephone; you must follow IBM's problem determination and resolution procedures. Scheduling of service will depend upon the time of your call and is subject to parts availability. Service levels are response time objectives and are not guaranteed. The specified level of warranty service may not be available in all worldwide locations; additional charges may apply outside IBM's normal service area. Contact your local IBM representative or your reseller for country and location specific information.

IBM On-Site Repair (IOR) IOR, 24 hours a day, 7 days a week, same-day response.

IBM will provide repair services for the failing machine at your location and verify its operation. You must provide suitable working area to allow disassembly and

reassembly of the IBM machine. The area must be clean, well lit, and suitable for the purpose. Depending on the proximity of the master console to the SAN File System cluster, you may also need to provide a keyboard, monitor, and mouse for attachment to a SAN File System engine.

Getting help online

Be sure to visit the support page for the SAN File System, complete with FAQs, parts information, technical hints and tips, technical publications, and downloadable files, if applicable. This page is at: www.ibm.com/storage/support.

Before you call for service

Some problems can be solved without outside assistance, by using the online help, by looking in the online or printed documentation that comes with the SAN File System, or by consulting the support Web page noted in Table 7 on page 151. Also, be sure to read the information in any README files and release notes that come with the SAN File System.

Getting help by telephone

With the original purchase of the SAN File System, you have access to extensive support coverage. During the product warranty period, you may call the IBM Support Center (1 800 426-7378 in the U.S.) for product assistance covered under the terms of the hardware IBM warranty or the software maintenance contract that comes with product purchase.

Please have the following information ready when you call:

- Machine type and model or the SAN File System software identifier. The software identifier can be either the product name (SAN File System) or the Product Identification (PID) number.
- Serial numbers of the SAN File System engines, or your proof of purchase
- Description of the problem
- Exact wording of any error messages
- Hardware and software configuration information

If possible, have access to your computer when you call.

In the U.S. and Canada, these services are available 24 hours a day, 7 days a week. In the U.K., these services are available Monday through Friday, from 9:00 a.m. to 6:00 p.m. In all other countries, contact your IBM reseller or IBM marketing representative.[1]

1. Response time will vary depending on the number and complexity of incoming calls.

Appendix A. Accessibility

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully.

Features:

These are the major accessibility features in SAN File System:

- You can use screen-reader software and a digital speech synthesizer to hear what is displayed on the screen.

  Note: The SAN File System Information Center and its related publications are accessibility-enabled for the IBM Home Page Reader.
- You can operate all features using the keyboard instead of the mouse.

Navigating by keyboard:

You can use keys or key combinations to perform operations and initiate many menu actions that can also be done through mouse actions. You can navigate the SAN File System console and help system from the keyboard by using the following key combinations:

- To traverse to the next link, button or topic, press Tab inside a frame (page).
- To expand or collapse a tree node, press Right or Left arrows, respectively.
- To move to the next topic node, press Down arrow or Tab.
- To move to the previous topic node, press Up arrow or Shift+Tab.
- To scroll all the way up or down, press Home or End, respectively.
- To go back, press Alt+Left arrow
- To go forward, press Alt+Right arrow.
- To go to the next frame, press Ctrl+Tab. There are quite a number of frames in the help system.
- To move to previous frame, press Shift+Ctrl+Tab.
- To print the current page or active frame, press Ctrl+P.

Appendix B. Cabling diagram

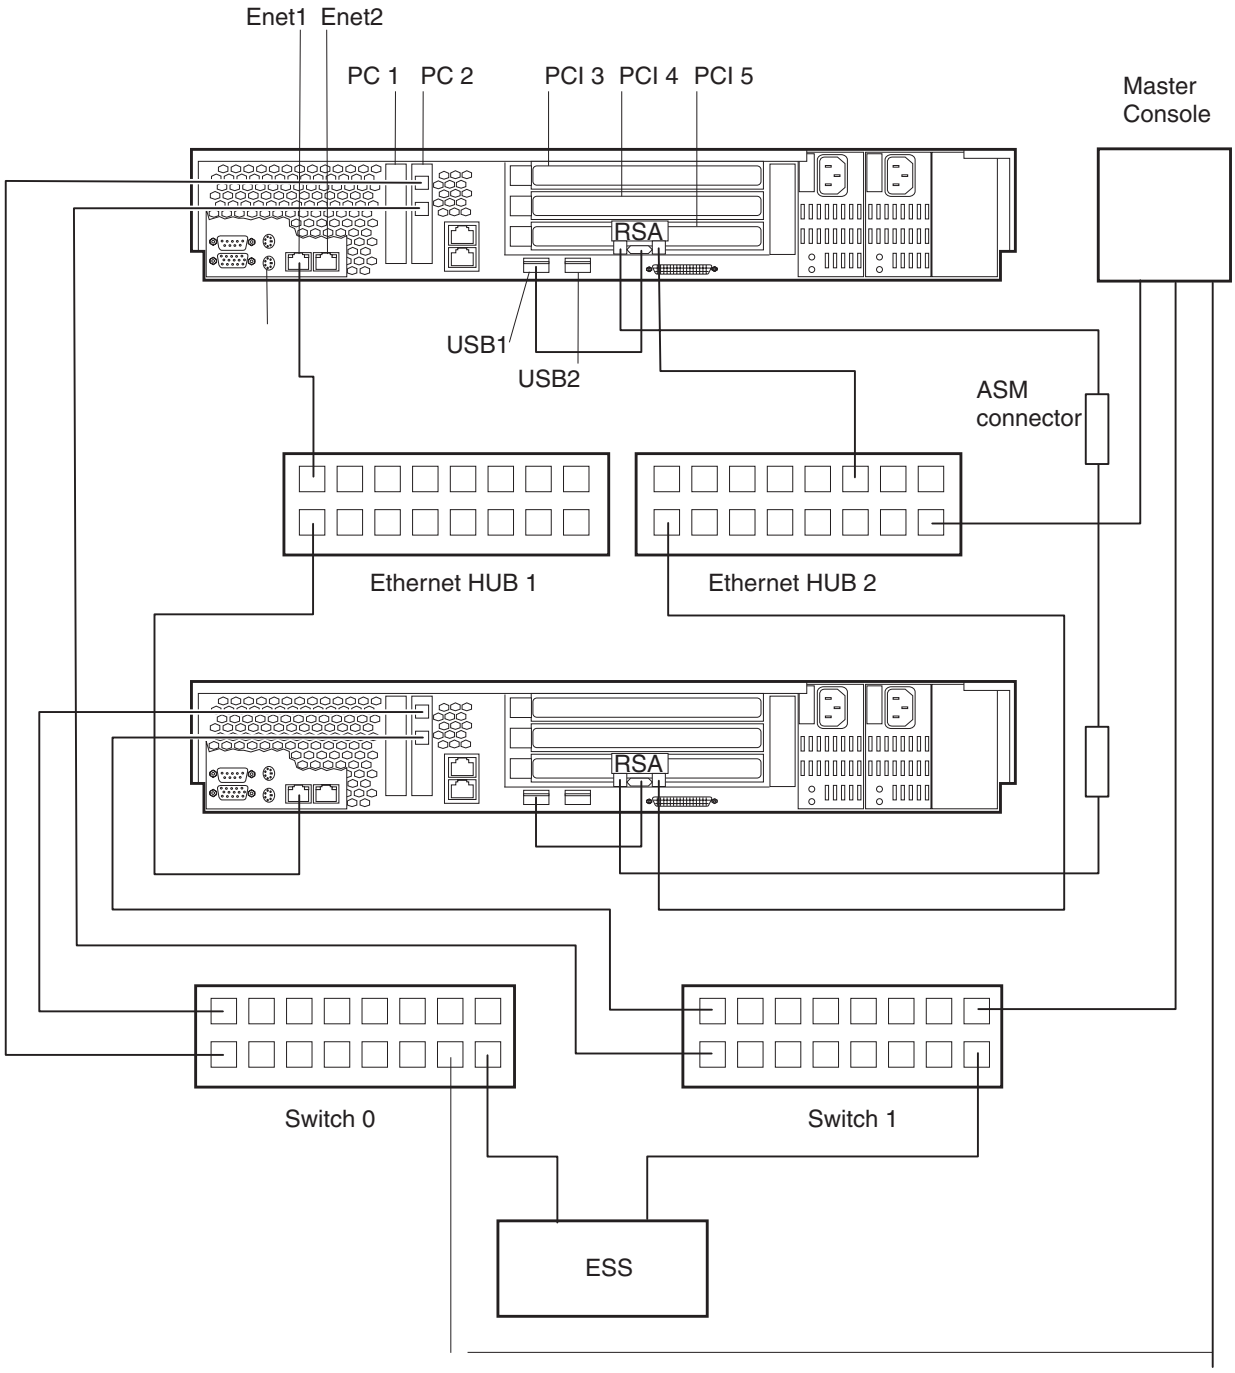

Figure 5. Two-engine, two-switch, two-hub cabling example

Appendix C. Commands

SAN File System has two sets of commands: administrative and client commands.

Administrative commands:

The administrative commands run on the storage engines that host the Metadata server. Most commands must be run from the master Metadata server. There are a few commands that must be run from subordinate Metadata server for specific situations.

You run a majority of the administrative commands from the tanktool session to manage SAN File System. There are a few commands that must be run from the operating-system shell prompt.

To use the administrative commands, you must log in directly to the engine, or from another workstation through SSH, using the local operating system authentication mechanism. You must then log in to the Administrative server on the engine using the same administrative user ID and password that you would use to log into the SAN File System console. You can specify the password in one of two ways:

- Set the password in the sclif.properties file, located in your home directory on the engine (for example, joe/sclif.properties), to your valid LDAP password using the **tankpasswd** utility.
- Set the SFS_CLI_PASSWDFILE environment variable to the location of the password file.

Note: The Administrative commands are case sensitive. If you enter a command in uppercase, you will receive an error.

Client commands:

The client commands runs on any client machine on which the client file-system driver has been installed. It provides a set of commands that you can use to manage your clients.

To use the client commands, you must log in directly to the client machine or from another workstation using SSH. You log in using the user ID and password for the client machine. You must have administrative (Windows) or root (UNIX-based) privileges to use the client commands.

Related topics:

- "Accessing an engine through SSH" on page 23
- "Administrative commands" on page 158
- "Client commands" on page 297

Administrative commands

The following table provides a brief description and administrative role for each command in the Administrative command-line interface.

Command	Description	Roles
General		
exit	Ends a tanktool session.	Monitor, Backup, Operator, Administrator
help	Displays a list of commands available in the Administrative command-line interface and optionally displays the syntax or brief description of each command.	Monitor, Backup, Operator, Administrator
quit	Ends the tanktool session.	Monitor, Backup, Operator, Administrator
setoutput	Sets the output format for the Administrative command-line interface.	Monitor, Backup, Operator, Administrator
Alerts and logs		
addsnmpmgr	Adds an SNMP manager to receive SNMP traps. To generate traps, you must add an SNMP manager and set the SNMP traps to be generated (using the **settrap**command).	Administrator
catlog	Displays the contents of the various log files maintained by the Administrative server and the cluster.	Monitor, Backup, Operator, Administrator
clearlog	Clears the audit log and cluster log files. Because the event log is a subset of the cluster log, the event log is also cleared when you clear the cluster log.	Administrator
lssnmpmgr	Displays a list of SNMP managers and their attributes.	Monitor, Backup, Operator, Administrator
lstrapsetting	Displays a list of event types that currently generate an SNMP trap.	Monitor, Backup, Operator, Administrator
rmsnmpmgr	Removes an SNMP manager (recipient).	Administrator
settrap	Specifies whether an SNMP trap is generated and sent to all SNMP managers when a specific type of event occurs on the Metadata server.	Administrator
Clients		
lsclient	Displays a list of clients that are currently being served by one or more Metadata servers in the cluster. You must be logged on to the master Metadata server to use this command.	Monitor, Backup, Operator, Administrator
Cluster		
chclusterconfig	Modifies the cluster settings that do not require a restart when changed.	Administrator

Command	Description	Roles
quiescecluster	Changes the state of all Metadata servers in the cluster to one of three quiescent states.	Operator, Administrator
resumecluster	Brings all Metadata servers in the cluster to the online state.	Operator, Administrator
startcluster	Starts all Metadata servers in the cluster and brings them to the full online state.	Operator, Administrator
statcluster	Displays status, network, workload, and configuration information about the cluster.	Monitor, Backup, Operator, Administrator
stopcluster	Stops all Metadata servers in the cluster gracefully.	Operator, Administrator
upgradecluster	Verifies that all the engines in the cluster have been upgraded to the new version of the software and then initiates the cluster upgrade process.	Administrator
Engines		
collectdiag	Collects data on all SAN File System components running in the system and saves the diagnostic results to various files on one or more specified engines.	Operator, Administrator
lsengine	Displays a list of storage engines and their attributes.	Monitor, Backup, Operator, Administrator
restartengine	Reboots one or more storage engines.	Operator, Administrator
startengine	Starts the CPU and operating system on one or more storage engines.	Operator, Administrator
statengine	Displays status information about a specific storage engine.	Monitor, Backup, Operator, Administrator
stopengine	Shuts down the operating system and powers off one or more storage engines.	Operator, Administrator
Filesets (containers)		
attachcontainer	Attaches an existing detached fileset (container) to a specific point in the global namespace or moves an already attached fileset to a new attach location.	Administrator
chcontainer	Modifies the settings of one or more filesets (containers).	Administrator
detachcontainer	Detaches one or more filesets (containers) from the global namespace.	Administrator
lscontainer	Displays a list of filesets (containers) and their attributes.	Monitor, Backup, Operator, Administrator
mkcontainer	Creates a new fileset (container).	Administrator
rmcontainer	Removes one or more empty, detached filesets (containers) and optionally the files in the filesets, including any FlashCopy images.	Administrator
setcontainerserver	Reassigns an existing fileset (container) to be hosted by a different Metadata server.	Administrator

Command	Description	Roles
statcontainer	Displays the number of started and completed transactions for the filesets (containers) being served by the local Metadata server.	—
FlashCopy images		
lsimage	Displays a list of FlashCopy images in all filesets or a specified fileset.	Monitor, Backup, Operator, Administrator
mkimage	Creates a near-instantaneous FlashCopy image of the file layout and contents of the specified fileset and stores it in that fileset.	Backup, Operator, Administrator
reverttoimage	Reverts the current fileset to a specified FlashCopy image of the file layout and contents.	Administrator
rmimage	Deletes one or more FlashCopy images for a specific fileset.	Backup, Operator, Administrator
Metadata		
builddrscript	Converts the system-metadata disaster-recovery file, created by the **mkdrfile** command, into a set of recovery scripts. You would then run these scripts to recreate the system metadata in the event of a disaster.	Backup, Operator, Administrator
lsdrfile	Displays a list of system-metadata disaster-recovery files.	Monitor, Backup, Operator, Administrator
mkdrfile	Creates a new system-metadata disaster-recovery dump file	Backup, Operator, Administrator
rmdrfile	Deletes an existing system-metadata disaster-recovery dump file.	Backup, Operator, Administrator
startmetadatacheck	Starts the utility that performs a consistency check on the metadata for the entire system or a set of filesets (containers), generates reports in the cluster log, and optionally repairs inconsistencies in the metadata.	Administrator
stopmetadatacheck	Stops the metadata check utility that is currently in progress.	Administrator
Metadata servers		
lsautorestart	Displays a list of Metadata servers and the automatic-restart settings for each.	Monitor, Backup, Operator, Administrator
lsserver	Displays a lists of all Metadata servers in the cluster and their attributes if issued from the master Metadata server, or displays attributes about the local Metadata server if issued from a subordinate Metadata server.	Monitor, Backup, Operator, Administrator
setmaster	Sets a subordinate Metadata server as the new master Metadata server in the event of an irrecoverable loss of the current master Metadata server.	Administrator
startautorestart	Enables the Metadata server to restart automatically if it is down.	Administrator

Command	Description	Roles
startserver	Starts the specified Metadata server.	Operator, Administrator
statserver	Displays status, configuration, and workload information for a specific Metadata server in the cluster, if issued from the master Metadata server. Lists status, configuration, and workload information for the local Metadata server, if issued from a subordinate Metadata server.	Monitor, Backup, Operator, Administrator
stopautorestart	Disables the Metadata server from restart automatically if it is down.	Administrator
stopserver	Shuts down a subordinate Metadata server gracefully.	Operator, Administrator
Policies		
catpolicy	Displays the rules of the specified policies.	Monitor, Backup, Operator, Administrator
lspolicy	Lists the active and inactive policies.	Monitor, Backup, Operator, Administrator
mkpolicy	Creates a policy.	Administrator
rmpolicy	Deletes one or more inactive policies.	Administrator
usepolicy	Directs the Metadata server to make an existing policy the active policy and applies its rules to all subsequent file creations.	Administrator
Processes		
lsproc	Displays a list of long-running processes that are not yet complete and their attributes.	Monitor, Backup, Operator, Administrator
Storage pools		
chpool	Modifies the settings of one or more storage pools.	Administrator
lspool	Displays a list of the existing storage pools and their attributes.	Monitor, Backup, Operator, Administrator
mkpool	Defines one or more new user storage pool.	Administrator
rmpool	Deletes one or more empty, unreferenced storage pools.	Administrator
setdefaultpool	Changes a user storage pool to the default storage pool, and changes the previous default storage pool to a regular, nondefault user storage pool.	Administrator
Users		
lsadmuser	Displays a list of administrative users and their attributes.	Monitor, Backup, Operator, Administrator
resetadmuser	Forces all administrative users to log in again.	Administrator
Volumes and LUNs		

Command	Description	Roles
activatevol	Activates a suspended volume so that the Metadata server can allocate new data on the volume.	Administrator
chvol	Modifies the settings of one or more volumes.	Administrator
lslun	Lists the LUNs that are accessible from SAN File System.	Monitor, Backup, Operator, Administrator
lsvol	Displays a list of available volumes.	Monitor, Backup, Operator, Administrator
mkvol	Adds one or more volumes to a storage pool.	Administrator
reportvolfiles	Displays a list of files (and their attributes) in the specified volume.	Backup, Operator, Administrator
rmvol	Removes one or more volumes from a storage pool and redistributes the contents to other volumes in the same storage pool.	Administrator
suspendvol	Suspends one or more volumes so that the Metadata server cannot allocate new data on the volumes.	Administrator

activatevol

Activates a suspended volume so that the Metadata server can allocate new data on the volume.

Syntax

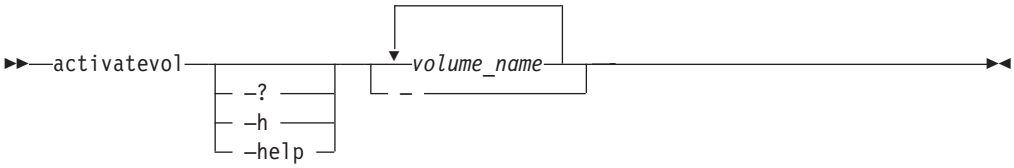

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

volume_name
Specifies the names of the one or more volumes to activate.

– Specifies that you want to read the names of one or more volumes to activate from stdin (for example, – << /work/vol_list.txt).

Description

Prerequisites:

1. You must have Administrator privileges to use the command.
2. You must be logged in to the engine that hosts the master Metadata server to run this command.

Examples

Activate two volumes The following example activates two volumes: vol1 and vol2.

```
tanktool> activatevol vol1 vol2
Volume vol1 activated.
Volume vol2 activated.
```

Related topics

- "chvol" on page 178
- "lsvol" on page 227
- "mkvol" on page 240
- "reportvolfiles" on page 245
- "suspendvol" on page 294

addsnmpmgr

Adds an SNMP manager to receive SNMP traps.

Syntax

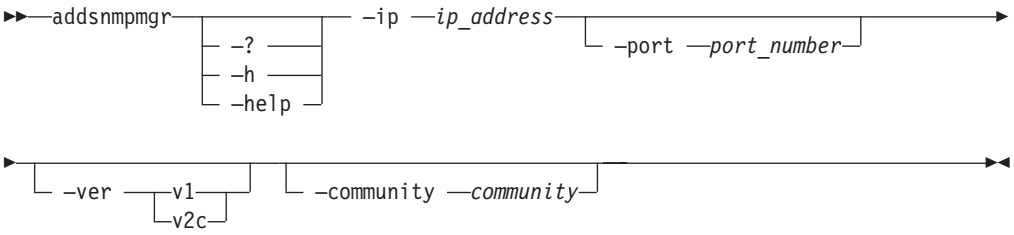

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–ip *ip_address*
Specifies the IP address of the SNMP trap receiver.

–port *port_number*
Specifies the SNMP-trap port number. The default port number is 162.

–ver v1 | v2c
Specifies the SNMP version. You can specify one of the following values:

v1 Sends SNMP version 1 traps. This is the default value.

v2c Sends SNMP version 2 notifications using community-based security.

–community *community*
Specifies the community as any alphanumeric string used to authenticate v1 or v2c traps. The default community is "public."

Commas and colons are not valid characters in the community name. This name can be no longer than 128 characters in length.

Description

Prerequisites:

1. You must have Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

Note: To generate traps, you must add an SNMP manager and set the SNMP traps to be generated (using the **settrap** command).

If an SNMP manager is not added, then traps are not generated.

You can add up to two SNMP managers to SAN File System. One of these SNMP managers may be used by the optional Call Home feature.

When you change a disruptive cluster setting, dynamic cluster settings (such as SNMP settings) cannot be modified until you reboot the cluster.

Examples

Adds an SNMP manager The following example adds an SNMP manager using SNMP v2c.

```
tanktool> addsnmpmgr –ip 192.168.0.1 –port 8192 –ver v2c –community SNMPMgr1
SNMP manager 192.168.0.1 successfully added.
```

Related topics

- "Service Alert" on page 16
- "catlog" on page 169
- "clearlog" on page 180
- "lssnmpmgr" on page 225
- "lstrapsetting" on page 226
- "rmsnmpmgr" on page 257
- "settrap" on page 268
- "SNMP" on page 18

attachcontainer

Attaches an existing detached fileset (container) to a specific point in the global namespace or moves an already attached fileset to a new attach location.

Syntax

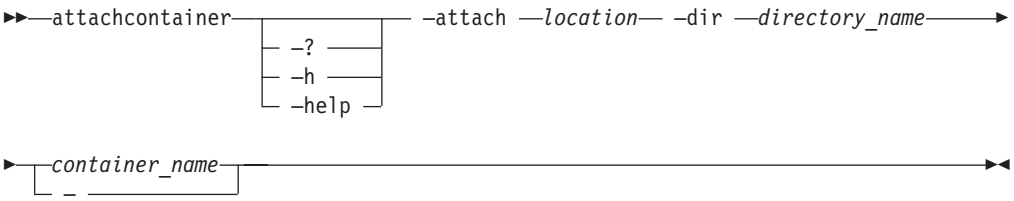

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–attach *location*
Specifies the directory path at which to attach the fileset. The root fileset in the global namespace must be included in the directory path. For example, to attach the fileset to the /home location, specify –attach sanfs/home. To attach the fileset directly to the root of the global namespace, use –attach sanfs.

Use only forward slashes (/) in the directory path for delimiters.

This directory ***must*** already exist.

Note: The name of the root fileset is the same as the name of the cluster. This name is "sanfs."

–dir *directory_name*
Specifies the directory name for the root of the fileset directory tree. The directory name ***must not*** already exist. This directory can be up to 256 characters in length.

The directory must not contain backslash (\) or forward slash (/) characters as delimiters.

The attach point (*location*/*directory_name*) appears as a drive and directory on clients for Windows or as a file system for client based on UNIX. It is recommended that you map the name of the directory to its associated fileset by using the same name as the fileset plus an additional suffix to distinguish it as the directory and not the fileset. For example, if you want to attach the fileset named work to the root of the global namespace, you would specify "sanfs" for the *location* and specify "work_ap" for the *directory_name*. On a client for Windows, the "work_ap" directory would appear under the "sanfs" drive.

container_name
Specifies the name of the fileset (or container) to attach.

– Specifies that you want to read the name of the fileset to attach from stdin (for example, – << /work/cnt_list.txt).

Description

Prerequisites:

1. You must have Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

You make a fileset part of the global namespace by attaching it to a specific location. To a client, this attach point looks like a regular directory.

Examples

Attaching a fileset The following example attaches fileset cntA to the sanfs/homes location in the global namespace using cntA_ap as its directory name.

```
tanktool> attachcontainer –attach sanfs/homes –dir cntA_ap cntA
Container cntA attached to directory sanfs/homes/cntA_ap.
```

Related topics

- “chcontainer” on page 174
- “detachcontainer” on page 183
- “lscontainer” on page 198
- “mkcontainer” on page 230
- “rmcontainer” on page 251

builddrscript

Converts the system-metadata disaster-recovery file into a set of recovery scripts used to recreate the system metadata in the event of a disaster.

Syntax

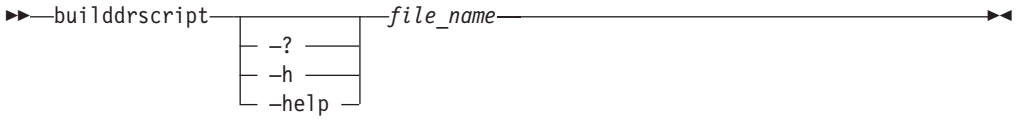

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

file_name
Specifies the name of an existing system-metadata disaster-recovery file.

Description

Prerequisites:: You must have Backup, Operator, or Administrator privileges to use the command.

Note: This command operates only against the local engine.

This command is used in conjunction with the **mkdrfile** command to create a set of recovery scripts that are used to recreate the system metadata in the event of a disaster. The **builddrscript** command is run against the system-metadata disaster-recovery file created by the **mkdrfile** command.

Note: The output for this command is written to the /usr/tank/server/DR directory. This command will overwrite any files that were created by a previous run of this command. If you want to preserve the existing files, copy them to another directory.

You would run this command against the system-metadata disaster-recovery file only if you are intending to recreate the system metadata.

After this command generates the recovery scripts, you must edit the scripts before running them.

Examples

Builds the recovery scripts The following example builds the disaster recovery scripts from the system-metadata disaster-recovery file named dr1.

```
tanktool> builddrscript dr1
Disaster recovery script files for "dr1" were built successfully.
```

Related topics

- “mkdrfile” on page 233
- “lsdrfile” on page 202
- “rmdrfile” on page 252

catlog

Displays the contents of the various log files maintained by the Administrative server and the cluster.

Syntax

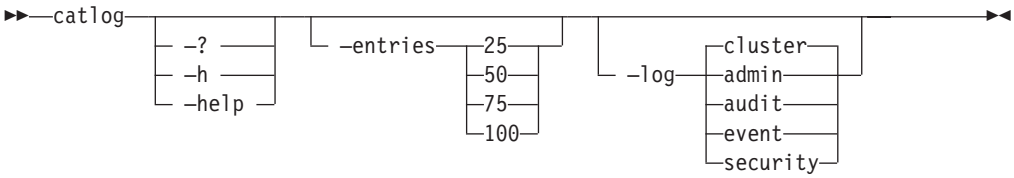

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–entries
Specifies the number of log entries to show at a time, from oldest to newest. Valid values are 25, 50, 75, or 100. If not specified, this command shows the entire log.

–log
Displays entries in the specified log, ordered by timestamp starting with the most recent entry. The default is cluster.

admin Displays entries in the administrative log, which maintains a history of messages created by the Administrative server.

audit Displays entries in the audit log, which maintains a history of all commands issued by any administrator for all Metadata servers in the cluster.

cluster
Displays entries in the cluster log, which maintains a history of messages created by all Metadata servers in the cluster.

event Displays entries in the event log, which maintains a history of event messages issued by all Metadata servers in the cluster.

security
Displays entries in the security log, which maintains a history of administrative-user login activity.

Description

Note: If you run this command from an engine hosting a subordinate Metadata server, logs for only the local engine are displayed. If you run this command from the engine hosting the master Metadata server, logs for the entire cluster are displayed.

If there are log entries that have not been displayed, you are prompted to press Enter to display the next set of entries or to type **exit** and press Enter to stop.

This command displays the following information for the specified log:

- Message code.

- Severity level (Severe, Error, Warning, Info).
- Message type (Normal or Event).
- Name of the Metadata server.
- Date and time the message was generated.
- Message.

Examples

Display the cluster log The following example displays the contents of the cluster log.

```
tanktool> catlog -log event

ID           Level   Server   Date and Time
=================================================
TANCM0383I   Info    ST1      Feb 20, 2003 8:39:15 PM
TANCM0384I   Info    ST2      Feb 20, 2003 8:40:46 PM
TANCM0384I   Info    ST2      Feb 21, 2003 1:05:27 AM
TANCM0383I   Info    ST1      Feb 21, 2003 1:05:18 AM

Message
=================================================
ALERT: The cluster's state has changed to Online(10)
ALERT: The server's state has changed to NotRunning(0)
ALERT: The server's state has changed to Online(10)
ALERT: The cluster's state has changed to Online(10)
```

Related topics

- "Administrative log" on page 37
- "Audit log" on page 38
- "clearlog" on page 180
- "Server log" on page 40
- "Security log" on page 39

catpolicy

Displays the rules of the specified policies.

Syntax

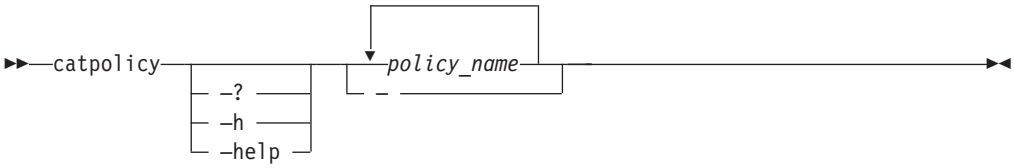

Description

Prerequisite: You must be logged in to the engine hosting the master Metadata server to run this command.

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

policy_name
Specifies the names of one or more policies to display.

– Specifies that you want to read the names of one or more policies to display from stdin (for example, – << /work/policies_list.txt).

Examples

Display the rules in a policy The following example displays the rules for the DEFAULT and the ActivePolicy policies:

```
tanktool> catpolicy  DEFAULT activePolicy
DEFAULT:
   VERSION 1
   rule 'stgRule1' set stgpool 'pool1' for CONTAINER ('cnt_A')
   rule 'stgRule2' set stgpool 'pool2' where NAME like '%.doc'
   rule 'stgRule3' set stgpool 'pool3' where DAYOFWEEK(CREATION_DATE) == 1
   rule 'stgRule4' set stgpool 'pool4' where USER_ID <= 100
```

Related topics

- “lspolicy” on page 213
- “mkpolicy” on page 236
- “rmpolicy” on page 255
- “usepolicy” on page 296

chclusterconfig

Modifies the cluster settings that do not require a restart when changed.

Syntax

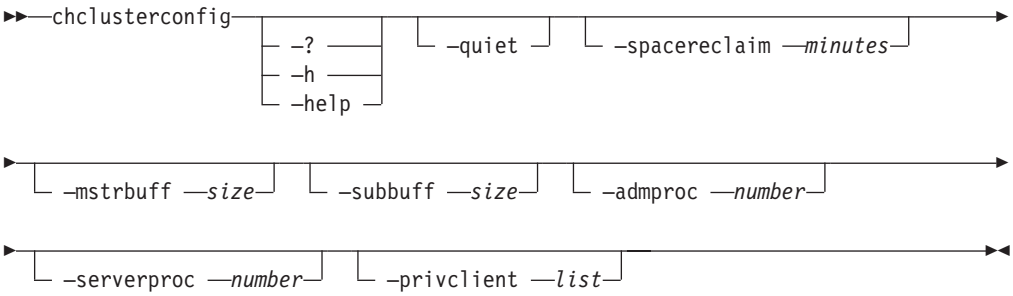

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–quiet
Turns off the prompt to confirm the change operation.

–spacereclaim *minutes*
Specifies the storage pool-space-reclamation interval, in minutes. You can specify a value from 0 to 1 440. The default is 60. A value of 0 disables storage pool-space reclamation.

–mstrbuff *size*
Specifies the size, in 4-KB pages, of the master Metadata server buffer. You can specify a value from 2 048 and 8 192. The default is 2 048.

Note: The buffer size can only be increased; it cannot be decreased.

Note: This parameter is intended for use only by trained service personnel.

–subbuff *size*
Specifies the size, in 4-KB pages, of the subordinate Metadata server buffer. You can specify a value from 30 000 and 250 000. The default is 200 000.

Note: The buffer size can only be increased; it cannot be decreased.

Note: This parameter is intended for use only by trained service personnel.

–admproc *number*
Specifies the maximum limit for administrative processes, including the number of Metadata server commands (from the administrative CLI or SAN File System console), that can run simultaneously. This limit is independent of the Server Workload Process Limit. A few commands, such as those that change the state of the cluster, can run on a temporary process even if the limit is currently exceeded. You can specify a value from 4 to 10. If not specified, the default value is 4.

Note: This limit can only be increased; it cannot be decreased.

Note: This parameter is intended for use only by trained service personnel.

–serverproc *number*
Specifies the maximum limit for the Metadata server workload process, which involves only Metadata server workload operations. You can specify a value from 10 to 50. The default value is 20.

Note: This limit can only be increased; it cannot be decreased.

Note: This parameter is intended for use only by trained service personnel.

–privclient *list*
Specifies a list of clients, separated by commas, that are to have root-level privileges (for example, –privclient saki,mefi,leki). The current list is replaced with the new list of clients. You can specify from zero to 30 clients; each client name can be no more than 256 characters in length.

Attention: This new list replaces the entire list of current privileged clients. If there are clients that have root or administrative privileges and you want to add one more client, you must specify both the current and new clients in the new list.

To remove all clients with root or administrative privileges, specify **–privclient** `""`; however, at least one privileged client is required to create files and directories under the root directory of filesets.

This command does not verify that the specified client names are valid because the Metadata server does not keep static information about the clients. It does not keep a list of allowed or possible clients. It has only a list of clients that it allows privileged access if they present themselves. This command lists the clients that are currently accessing data.

Description

Prerequisites:: You must have Administrator privileges to use the command.

If you modify any disruptive settings using the **chadvclusterconfig** command, you cannot modify any cluster settings until you restart the cluster.

Examples

Changes the cluster settings The following example changes space-reclamation setting for the cluster:

```
tanktool> chclusterconfig -spacereclaim 100
Are you sure you want to change cluster configuration settings? [y/n] y
Cluster successfully modified.
```

Related topics

- “chadvclusterconfig” on page 323
- “startcluster” on page 272
- “statcluster” on page 278
- “stopcluster” on page 290
- “upgradecluster” on page 295

chcontainer

Modifies the settings of one or more filesets (containers).

Syntax

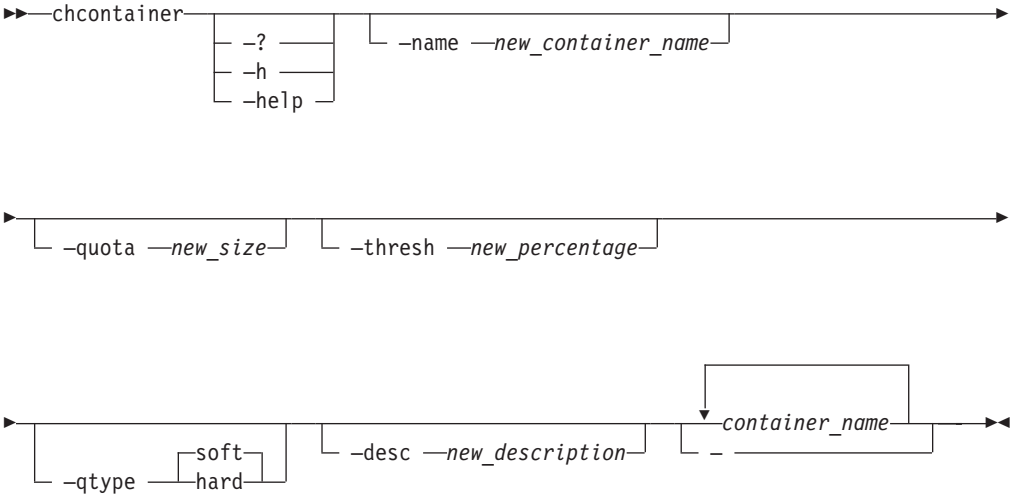

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–name *new_container_name*
Specifies the new name to assign to the fileset (or container). This name can be up to 256 characters in length.

–quota *new_size*
Specifies the new maximum size limit, in megabytes, for the specified fileset that, when exceeded, causes the Metadata server to generate an alert. You can specify a number from 0 to 1 073 741 824 MB (1 PB). If set to 0, there is no quota size limit for this fileset, and alerts are not sent.

Note: If a hard quota already exists, you must not set the new size to a value less than the current quota size for this fileset.

–thresh *new_percentage*
Specifies the new maximum percentage (alert threshold) of the specified quota size for this fileset that, when exceeded, causes the Metadata server to generate an alert. You can specify a value between 0 and 100. If set to 0, no alerts are generated.

–qtype hard | soft
Specifies the quota type for the fileset. You can specify one of the following values:

hard A hard quota produces a log message and potential alert when the quota is met, and denies requests for additional space.

soft A soft quota produces a log message and potential alert when the quota size is exceeded, but grants requests space allocation.

–desc *description*
Specifies a new description for the fileset. The description must be enclosed in matching single (') or double (") quotation marks if it contains any blank characters. This description can be up to 256 characters in length.

The default is an empty string.

container_name
Specifies the names of one or more filesets (or containers) to modify.

Note: If you specify a name using the **–name** option, you may specify one only fileset.

– Specifies that you want to read the names of one or more filesets to modify from stdin (for example, – << /work/cnt_list.txt).

Description

Prerequisites:

1. You must have Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

You cannot modify the name for multiple filesets at one time.

Examples

Modify parameters of two filesets The following example changes the quota and threshold values for two filesets (*cnt_A* and *cnt_B*). It, sets the quota to 1 000 MB and specifies to send an alert when the quota reaches 70%.

```
tanktool> chcontainer –quota 1000 –thresh 70 –qtype soft cnt_A cnt_B
Container cnt_A modified.
Container cnt_B modified.
```

Related topics

- "attachcontainer" on page 166
- "detachcontainer" on page 183
- "lscontainer" on page 198
- "mkcontainer" on page 230
- "rmcontainer" on page 251

chpool

Modifies the settings of one or more storage pools.

Syntax

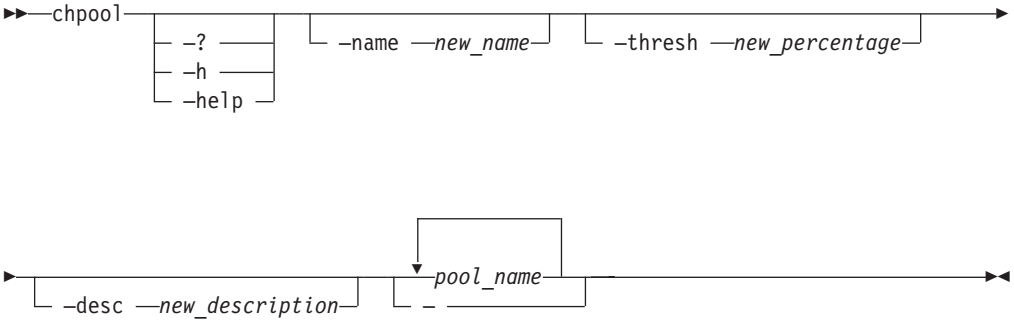

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–name *new_name*
Specifies the new name to assign to the storage pool. This name can be up to 256 characters in length.

Note: You cannot change the name of the SYSTEM storage pool.

–thresh *new_percentage*
Specifies a percentage of the storage pool's estimated capacity that, when reached or exceeded, causes the Metadata server to generate an alert. If the capacity drops and then reaches the percentage again, additional alerts are sent. You can specify a value between 0 and 100. If set to 0, no alert is generated.

–desc *new_description*
Specifies a description for the storage pool. The description must be enclosed in matching single (') or double (") quotation marks if it contains any blank characters. This description can be up to 256 characters in length.

The default is an empty string.

Note: You cannot change the description of the SYSTEM storage pool.

pool_name
Specifies the names of one or more storage pools to modify.

Note: If you specify a name using the **–name** option, you may specify only one storage pool.

– Specifies that you want to read the names of one or more storage pools to modify from stdin (for example, – << /work/stgpool_list.txt).

Description

Prerequisites:

1. You must have Administrator privileges to use the command.

2. You must be logged in to the engine hosting the master Metadata server to run this command.

You cannot modify the names of multiple storage pools at one time.

Examples

Modify settings for two storage pools The following example modifies the threshold setting for two storage pools: stgpool1 and stgpool2.

```
tanktool> chpool –thresh 70 stgpool1 stgpool2
Storage pool stgpool1 modified.
Storage pool stgpool2 modified.
```

Related topics

- “lspool” on page 216
- “mkpool” on page 238
- “rmpool” on page 256
- “setdefaultpool” on page 262

chvol

Modifies the settings of one or more volumes.

Syntax

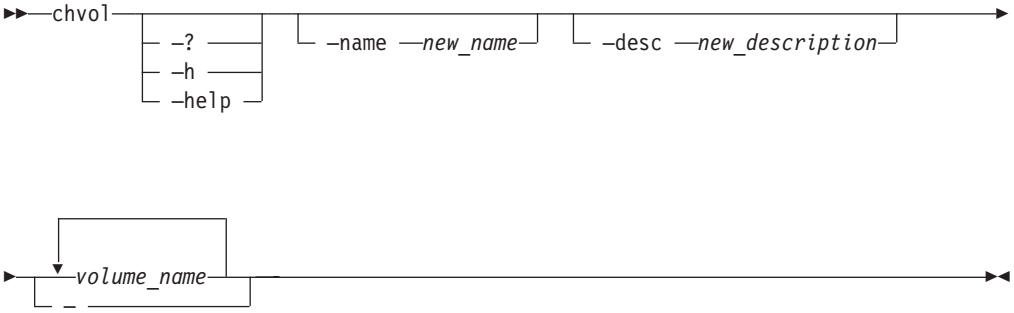

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–name *new_name*
Specifies the new name to assign to the volume. The volume name must be unique within a cluster and can be up to 256 characters in length.

–desc *new_description*
Provides a new description of the volume. The description must be enclosed in matching single (') or double (") quotation marks if it contains any blank characters and can be up to 256 characters in length.

The default is an empty string.

volume_name
Specifies the name of one or more volumes to modify.

Note: If you specify a name using the **–name** option, you may specify one only volume.

– Specifies that you want to read the names of one or more volumes to modify from stdin (for example, – << /work/vol_list.txt).

Description

Prerequisites:

1. You must have Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

You cannot modify the name of multiple volumes at one time.

Examples

Modify settings for multiple volumes The following example modifies the description of all the volumes listed in the file */tmp/vol_list.txt* to be "Engineering Volumes":

```
tanktool> chvol –desc "My SAN File System volumes" – << /tmp/vol_list.txt

Volume vol1 modified.
Volume vol2 modified.
Volume vol3 modified.
```

Related topics

- “activatevol” on page 163
- “lsvol” on page 227
- “mkvol” on page 240
- “reportvolfiles” on page 245
- “rmvol” on page 258
- “suspendvol” on page 294

clearlog

Clears the audit log and cluster log files.

Syntax

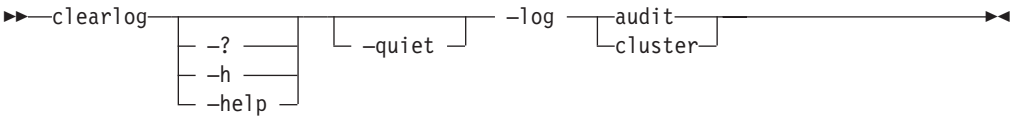

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–quiet
Turns off the prompt to confirm the clear operation

–log
Specifies the log file to clear.

audit Clears all entries in the audit log, which maintains a history of messages created by the Administrative server.

cluster
Clears all entries in the cluster log, which maintains a history of all messages issued by any administrator for all Metadata servers in the cluster.

Description

Prerequisites: You must have Administrator privileges to use the command.

Notes:

1. If you run this command from an engine hosting a subordinate Metadata server, the audit log and cluster log on only the local engine are cleared. If you run this command from the engine hosting the master Metadata server, the audit logs and cluster logs on all engines are cleared.
2. Because the event log is a subset of the cluster log, the event log is also cleared when you clear the cluster log.

Examples

Clear the cluster log The following example clears the cluster log entries.

```
tanktool> clearlog -log cluster
Are you sure you want to remove all entries in the cluster log? [y/n] y
Cluster log cleared.
```

Related topics

- “Audit log” on page 38
- “catlog” on page 169
- “Server log” on page 40

collectdiag

Collects data on all SAN File System components running in the system and saves the diagnostic results to various files on the local file system of each specified engine. This command is used to assist in problem determination.

Syntax

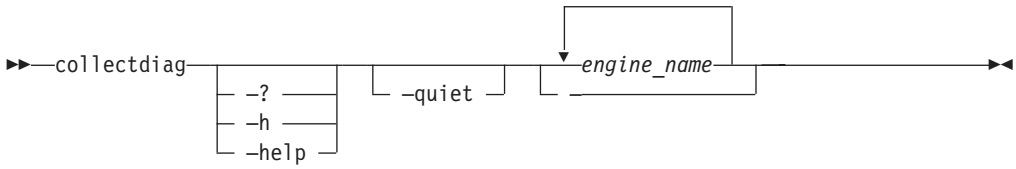

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–quiet
Turns off the prompt to confirm this operation

engine_name
Specifies the names of one or more engines for which to collect diagnostic data.

Note: You must be logged on to the master Metadata server to issue this command against one or more non-local engine. If you are not logged on to the master Metadata server, this command collects diagnostics data only for the local engine.

– Specifies that you want to read the name of the engine for which to gather diagnostic data from stdin (for example, – << /work/vol_list.txt).

Description

Prerequisites: You must have Operator or Administrator privileges to use the command.

Diagnostic data is saved on each specified engine in a subdirectory of the directory specified by the TANKDIR environment variable. The default directory is /usr/tank/pmf. The subdirectory name is the timestamp when the data is collected (for example. ,/usr/tank/pmf/20030507_081010).

Stdout and stderr output that is created by this command is saved in the /tmp/odbcoutfile on the local disk of each specified engine.

You must be logged on to the master Metadata server to issue this command against a non-local engine. Otherwise, this command works only on the local engine.

Examples

Collect diagnostic data The following example collects diagnostic data for engine ST1:

```
tanktool> collectdiag ST1
Are you sure you want to collect diagnostic data for ST1 since this
procedure may take a few minutes to complete? [y/n]  y
Diagnostic file successfully created for Engine ST1. Tip: See
usr/tank/pmf to read the diagnostic files.
```

Related topics

- "Engines" on page 5

detachcontainer

Detaches one or more filesets (containers) from the global namespace.

Syntax

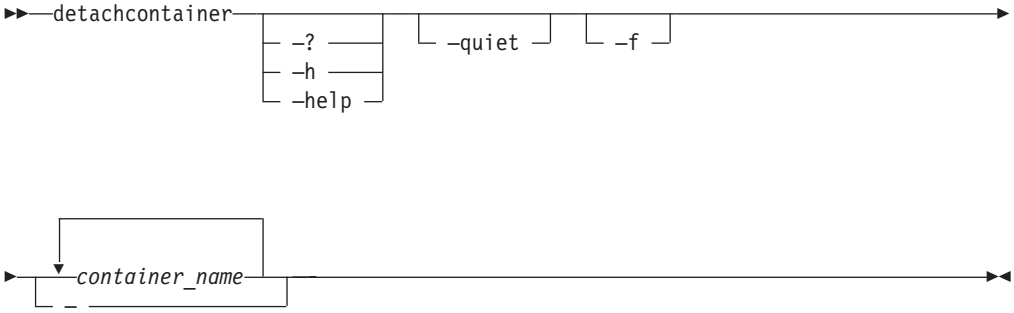

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–quiet
Turns off the prompt to confirm the detach operation.

–f Forces the Metadata server to continue the detach operation even if there are open files in the fileset and suppresses any warning messages.

container_name
Specifies the names of one or more filesets (or containers) to detach.

– Specifies that you want to read the names of one or more filesets to detach from stdin (for example, – << /work/cnt_list.txt).

Description

Prerequisites:

1. You must have Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

The specified fileset ***must not*** have nested filesets. You must detach any nested filesets before detaching the fileset.

Unless you specify the **–f** parameter, there must not be any open files or any locks on files in the fileset.

The directory to which the fileset was attached is deleted after this operation is complete.

Examples

Detaching a fileset The following example detaches the fileset named *cnt_A*.

```
tanhktool> detachcontainer cnt_A
Are you sure you want to detach container cnt_A? [y/n] y
Container cnt_A  detached from directory sanfs/homes.
```

Related topics

- "attachcontainer" on page 166
- "chcontainer" on page 174
- "lscontainer" on page 198
- "mkcontainer" on page 230
- "rmcontainer" on page 251

exit

Ends a tanktool session.

Syntax

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–l Displays a list of available commands with the syntax diagrams for each.

Examples

Ends a tanktool session The following example ends the tanktool session:

```
tanktool> exit
shell>
```

help

Displays a list of commands available in the Administrative CLI and optionally displays the syntax or brief description of each command.

Syntax

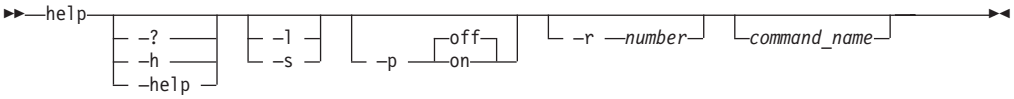

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–l Displays a list of available commands with the syntax diagrams for each. If you specify a command name with this parameter, this command displays the syntax for only the specified command.

–s Displays a list of available commands with a brief description of each. If you specify a command name with this parameter, this command displays a brief description for only the specified command.

–p Specifies whether to display one page of text at a time or all text at once.

off Displays all text at one time. This is the default value.

on Displays one page of text at time. Pressing any key displays the next page.

–r *number*
Specifies the number of rows per page to display when the **–p** parameter is on. The default is 24 rows. You can specify a value from 1 to 100.

command_name
Displays help information for the specified command, including the syntax diagram, parameter descriptions, return codes and errors, descriptions, examples, and miscellaneous remarks.

Description

If you specify this command with no parameters, this command displays only a list of available commands.

Examples

Display a description of a command The following example displays the description of the **mkimage** command:

```
tanktool>help -s lspool
lspool    Displays a list of existing storage pools and their attributes.
```

lsadmuser

Displays a list of administrative users and their attributes.

Syntax

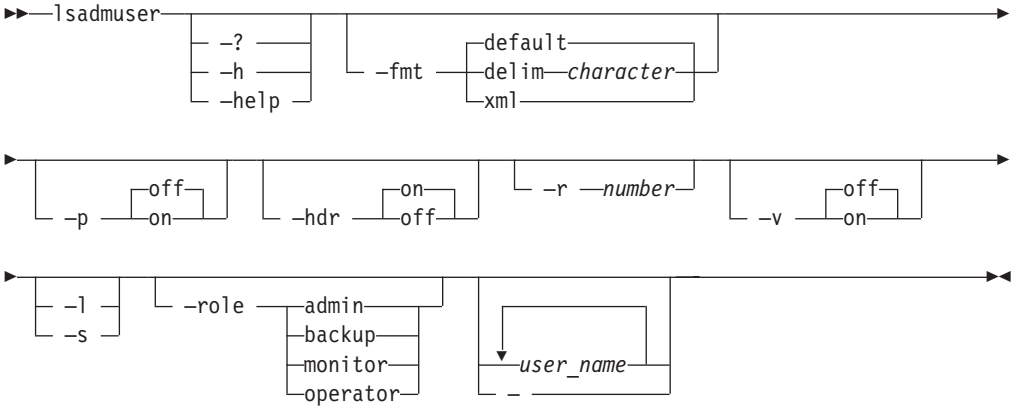

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–fmt
Specifies the format of the output. You can specify one of the following values:

default
Specifies to display output in a tabular format using spaces as the delimiter between the columns. This is the default value. For example:

```
Name   User Role  Authorization    Authorization Timeout (secs)
================================================================
maki   Admin      Current                                 300
```

delim *character*
Specifies to display output in a tabular format using the specified character to separate the columns. If you use a shell metacharacter (for example, * or \t) as the delimiting character, enclose the character in single quotation marks (') or double quotation marks ("). A blank space is not a valid character. For example:

```
maki,Admin,Current,300 secs
```

xml Specifies to display output using XML format, for example:

```
<IRETURNVALUE>
<INSTANCE CLASSNAME="STC_AdminUser">
<PROPERTY NAME="Name" TYPE="string"><VALUE TYPE='string'>maki</VALUE>
</PROPERTY>
<PROPERTY NAME="EffectiveRole" TYPE="unit16"><VALUE="unit16">3</VALUE>
</PROPERTY>
<PROPERTY NAME="IsAuthorizationCurrent" TYPE="boolean"><VALUE
TYPE="boolean">true</VALUE></PROPERTY>
<PROPERTY NAME="AuthCurrentRemainingTime" TYPE="uint32"><VALUE
TYPE="uint32">300</VALUE></PROPERTY>
</INSTANCE>
</IRETURNVALUE>
```

–p Specifies whether to display one page of text at a time or all text at once.

off Displays all text at one time. This is the default value when the **tanktool** command is run in single-shot mode.

on Displays one page of text at time. Pressing any key displays the next page. This is the default value when the **tanktool** command is run in interactive mode.

–hdr
Specifies whether to display the table header.

on Displays the table header. This is the default value.

off Does not display the table header.

–r *number*
Specifies the number of rows per page to display when the **–p** parameter is on. The default is 24 rows. You can specify a value from 1 to 100.

–v Specifies whether to enable verbose mode.

off Disables verbose mode. This is the default value.

on Enables verbose mode.

–l Displays additional information about the administrative users.

–s Displays only the names of the administrative users.

–role
Displays information for only those users with the specified user role . Possible roles are:

admin Shows only those users assigned to the Administrator role.

backup
Shows only those users assigned to the Backup role.

monitor
Shows only those users assigned to the Monitor role.

operator
Shows only those users assigned to the Operator role.

user_name
Specifies the names of one or more administrative users to list. If not specified, this command lists all of the administrative users assigned to all user roles.

– << *file_name*
Specifies that you want to read the names of the one or more administrative users to list from stdin (for example, – << /work/users_list.txt).

Description

Note: This command operates only against the local engine.

Administrative users are defined in the LDAP server as authorized to perform administrative actions on the SAN File System. A user's role is determined by its membership in an LDAP group that is associated with the various SAN File System roles.

If you do not specify a listing parameter, this command displays the following default information for each administrative user:

- Name of the administrative user.
- Role assigned to the administrative user.

- Authorization (Current or Not Current).

If you specify the **–l** listing parameter, this command displays the following information in addition to the default information for each administrative user:

- Authorization timeout, in seconds.

Examples

Lists all administrative users The following example lists all SAN File System administrative users and additional information about their assigned roles.

```
tanktool> lsadmuser –l

Name   User Role  Authorization    Authorization Timeout (secs)
============================================================
maki   Admin      Current                               300
saki   Backup     Not Current                             0
mizi   Operator   Current                               465
jeko   Monitor    Not Current                             0
```

Related topics

- "resetadmuser" on page 246

lsautorestart

Displays a list of Metadata servers and the automatic-restart settings for each.

Syntax

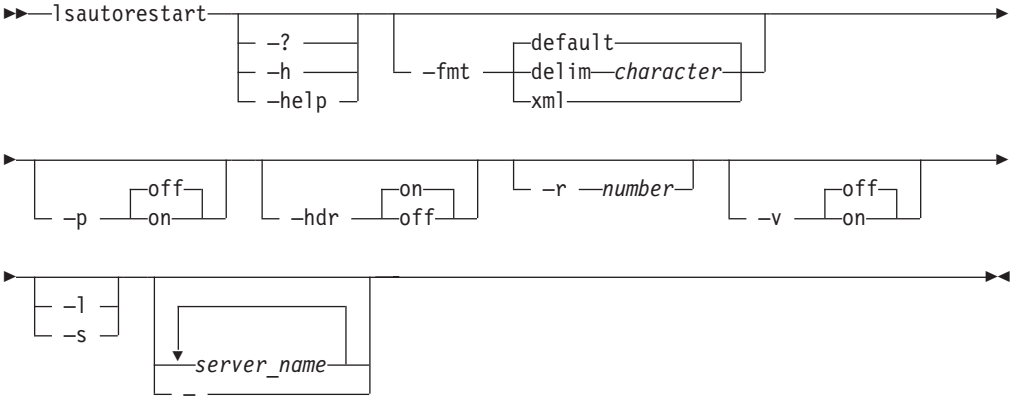

Parameters

–? | –h | –help

Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–fmt

Specifies the format of the output. You can specify one of the following values:

default

Specifies to display output in a tabular format using spaces as the delimiter between the columns. This is the default value. For example:

```
Name   Service State   Probe State   Last Probe                Probes
=========================================================================
ST1    Running         Live Server   Feb 11, 2003 9:35:47 AM   218690

Highest Retries
===============
              2
```

delim *character*

Specifies to display output in a tabular format using the specified character to separate the columns. If you use a shell metacharacter (for example, * or \t) as the delimiting character, enclose the character in single quotation marks (') or double quotation marks ("). A blank space is not a valid character. For example:

```
ST1,Running,Live Server,Feb 11, 2003 9:35:47 AM,218690,2
```

xml Specifies to display output using XML format, for example:

```
<IRETURNVALUE>
<INSTANCE CLASSNAME="STC_TankWatchdog">
<PROPERTY NAME="ServiceName" TYPE="string"><VALUE TYPE="string">ST1
</VALUE></PROPERTY>
<PROPERTY NAME="State" TYPE="uint32"><VALUE TYPE="uint32">1</VALUE>
</PROPERTY>
<PROPERTY NAME="ProbeState" TYPE="uint32"><VALUE TYPE="uint32">2</VALUE>
</PROPERTY>
<PROPERTY NAME="LastProbeTimeStamp" TYPE="datetime"><VALUE
TYPE="datetime">20030211093547&#46;000009&#45;420</VALUE></PROPERTY>
<PROPERTY NAME="TotalProbes" TYPE="uint64"><VALUE TYPE="uint64">218690
```

```
</VALUE></PROPERTY>
<PROPERTY NAME="RetriesHWM" TYPE="uint32"><VALUE TYPE="uint32">2</VALUE>
</PROPERTY>
</INSTANCE>
</IRETURNVALUE>
```

–p Specifies whether to display one page of text at a time or all text at once.

off Displays all text at one time. This is the default value when the **tanktool** command is run in single-shot mode.

on Displays one page of text at time. Pressing any key displays the next page. This is the default value when the **tanktool** command is run in interactive mode.

–hdr
Specifies whether to display the table header.

on Displays the table header. This is the default value.

off Does not display the table header.

–r *number*
Specifies the number of rows per page to display when the **–p** parameter is on. The default is 24 rows. You can specify a value from 1 to 100.

–v Specifies whether to enable verbose mode.

off Disables verbose mode. This is the default value.

on Enables verbose mode.

–l Displays additional information about the Metadata servers.

–s Displays default information about the Metadata servers.

server_name
Specifies the names of one or more Metadata servers to display. If not specified, settings for all Metadata servers are displayed.

– Specifies that you want to read the names of one or more Metadata servers to display from stdin (for example, – << /work/server_list.txt).

Description

Note: If you run this command from an engine hosting a subordinate Metadata server, you can display information about only the local Metadata server. If you run this command from the engine hosting the master Metadata server, you can display information about any Metadata server.

If you do not specify a listing parameter, this command displays the following default information for each Metadata server:

- Name of the Metadata server.
- State of the Metadata server restart service. Possible states are:
  Off The Metadata server restart service is not running.
  Failed The Metadata server restart service is running but has failed to restart the Metadata server after several attempts.
  Running
  The Metadata server restart service is running and ready to restart the Metadata server if it fails.
  Standby
  The Metadata server restart service is running, but the Metadata server has been manually shut down. You must manually restart Metadata server to return the Metadata server restart service to the running state.

Unknown
The Metadata server restart service is in an unknown state because the Metadata server could not be reached.

- Probed Metadata server status. Possible status values are:

  Absent Server
  The Metadata server restart service could not find the Metadata server and will attempt to restart the Metadata server.

  Live Server
  The Metadata server restart service found the Metadata server and it is up and running. There is no need to restart the Metadata server.

  Not Probed
  The Metadata server restart service has not started probing the Metadata server because it is either in the off or standby state.

  Probing
  The Metadata server restart service has started probing the Metadata servers for status.

  Unknown
  The Metadata server restart service cannot determine whether the Metadata server is live or absent.

  It is possible to for the probe state to be unknown when the Metadata server is live. For example, if the Metadata server fails, the Metadata server restart service determines that the Metadata server is absent, and restarts it. If the Metadata server is slow to start, on the next probe cycle, the Metadata server is neither live nor absent. When it is probed again, the Metadata server is live.

  Another example is if the local disk is full, and the Metadata server restart service tries to restart the Metadata server but cannot write to its log file. The Metadata server crashes, and the Metadata server restart service tries to restart it again. On the next probe cycle, the state of the Metadata server is absent or unknown.
- Timestamp when the Metadata server was last probed for status.
- Total number of probes that the Metadata server restart service performed.
- Highest number of times that the Metadata server restart service has attempted to restart the Metadata server.

If you specify the **–l** listing parameter, this command displays the following information in addition to the default information for each Metadata server:

- Timestamp when the automatic-restart service started on the Metadata server.
- Amount of time, in milliseconds, between each probe.
- Current number of times that the automatic-restart service has attempted to restart the Metadata server.
- Total number of times that the automatic-restart service has attempted to restart the Metadata server.
- Maximum number of times that the automatic-restart service will attempt to restart the Metadata server.
- Lowest number of times that the automatic-restart service has attempted to restart the Metadata server.
- Total number of times that a probe for a live Metadata server has taken longer than the test timeout interval.
- Maximum amount of time, in milliseconds to wait for a Metadata server to respond before declaring it down.
- Amount of time, in milliseconds, used the last time the Metadata server was probed to see if it is running.
- Highest amount of time, in milliseconds, taken to determine if the Metadata server is running.

- Lowest amount of time, in milliseconds, taken to determine if the Metadata server is running.
- Number of times the watchdog probed the Metadata server to determine if it is absent.
- Amount of time, in milliseconds, used the last time the Metadata server was probed to determine if it is absent.
- Highest amount of time, in milliseconds, taken to determine if the Metadata server is absent.
- Lowest amount of time, in milliseconds, taken to determine if the Metadata server is absent.

Examples

Displays the automatic-restart settings The following example displays the automatic-restart settings and additional information for all Metadata servers in the cluster.

```
tanktool> lsautorestart ST1

Name   Service State   Probe State   Last Probe                Probes
=====================================================================
ST1    Running         Live Server   Feb 11, 2003 9:35:47 AM        5

Highest Retries
===============
              2
```

Related topics

- "Metadata server" on page 8
- "startautorestart" on page 270
- "stopautorestart" on page 289

lsclient

Displays a list of clients that are currently being served by one or more Metadata servers in the cluster.

Syntax

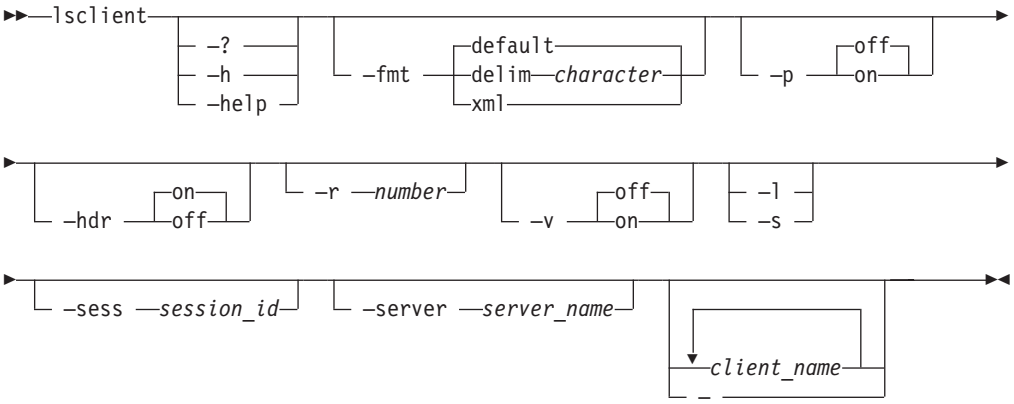

Parameters

–? | –h | –help

Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–fmt

Specifies the format of the output. You can specify one of the following values:

default

Specifies to display output in a tabular format using spaces as the delimiter between the columns. This is the default value. For example:

```
Client    Session ID   State     Server   Renewals   Privilege
============================================================
jeko               1   Current   ST1           45   Standard
```

delim *character*

Specifies to display output in a tabular format using the specified character to separate the columns. If you use a shell metacharacter (for example, * or \t) as the delimiting character, enclose the character in single quotation marks (') or double quotation marks (″). A blank space is not a valid character. For example:

```
jeko,1,Current,ST1,45,Standard
```

xml Specifies to display output using XML format, for example:

```
<IRETURNVALUE>
<INSTANCE CLASSNAME="STC_RegisteredFSClient">
<PROPERTY NAME="Name" TYPE="string"><VALUE TYPE="string">jeko</VALUE>
</PROPERTY>
<PROPERTY NAME="Id" TYPE="unit64"><VALUE TYPE="unit64">1</VALUE>
</PROPERTY>
<PROPERTY NAME="State" TYPE="unit16"><VALUE TYPE="unit16">1</VALUE>
</PROPERTY>
<PROPERTY NAME="SystemName" TYPE="string"><VALUE TYPE="string">ST1
</VALUE></PROPERTY>
<PROPERTY NAME="LeaseRenewals" TYPE="uint64"><VALUE TYPE="uint64">45
```

```
</VALUE></PROPERTY>
<PROPERTY NAME="IsPrivilege" TYPE="boolean"><VALUE TYPE="boolean">true
</VALUE></PROPERTY>
</INSTANCE>
</IRETURNVALUE>
```

–p Specifies whether to display one page of text at a time or all text at once.

off Displays all text at one time. This is the default value when the **tanktool** command is run in single-shot mode.

on Displays one page of text at time. Pressing any key displays the next page. This is the default value when the **tanktool** command is run in interactive mode.

–hdr
Specifies whether to display the table header.

on Displays the table header. This is the default value.

off Does not display the table header.

–r *number*
Specifies the number of rows per page to display when the **–p** parameter is on. The default is 24 rows. You can specify a value from 1 to 100.

–v Specifies whether to enable verbose mode.

off Disables verbose mode. This is the default value.

on Enables verbose mode.

–l Displays additional information about each client.

–s Displays limited information about each client.

–sess *session_id*
Specifies the client-session ID to list. If not specified, this command displays a list of all clients.

–server *server_name*
Specifies the name of a Metadata server to query for clients. If specified, this command displays only those clients that are served by the specified Metadata server. If not specified, this command displays clients that are served by all Metadata servers in the cluster.

client_name
Specifies the names of one or more clients to list. If not specified, this command displays all clients that are being served by the specified Metadata server.

– Specifies that you want to read the names of one or more clients to list from stdin (for example, – << /work/client_list.txt).

Description

Prerequisite: You must be logged on to the master Metadata server to use the command.

When you specify multiple parameters in this command, these behaviors apply:

- Flags work as AND coordinators.
- Objects work as OR coordinators.

- When flags and objects are combined, the they are listed in this manner: WHERE flag_1 AND flag_2 AND ... AND flag_n AND (object_1 OR object_2 OR ... OR object_m).
- When flags or objects are not required and not specified, this command searches all possible values unless a default overrides this rule.

For example, if you specify –sess, –server, and two client names, the search would be performed in the following manner: WHERE session_ID AND server_name AND (client_name_1 OR client_name_2).

If you do not specify a listing parameter, this command displays the following default information for each client:

- Client name. This name is configured by the administrator on the client node.
- Client-session ID. Each client can have one session running on each Metadata server simultaneously.
- State of the client. Possible states are:

  Current
  The client is currently being served by one of the Metadata servers.

  Expired
  The client is not currently being served by one of the Metadata servers.
- Metadata server that is serving the client.
- Number of renewals since the client-session was started.
- Client privilege or access level. Possible privileges are:

  Root Indicates that the user has root privileges, and is considered a privileged user.

  Standard
  Indicates that the user has default access level.

If you specify the **–s** listing parameter, this command displays only the following information for each client:

- Client name.
- Client-session ID.
- Metadata server that is serving the client.

If you specify the **–l** listing parameter, this command displays the following information in addition to the default information for each client:

- Date and time of the last client-session renewal.
- Time until the next client-session renewal is to occur, in seconds.
- Client IP address.
- Client port number.
- Client operating-system platform.
- Client file-system driver version.
- Number of transactions started, per client session.
- Number of transactions that have completed, per client session.
- Number of session locks, per client session.
- Number of data locks, per client session.
- Number of byte-range locks, per client session.

Note: The Metadata server does not know which clients are connected to SAN File System. If you specify a client that is not connected, this command will return a warning that no rows were found.

Examples

List clients The following example displays information about all clients currently being served by Metadata server ST1.

```
tanktool> lsclient -server ST0,ST1

Client    Session ID   State     Server   Renewals   Privilege
=============================================================
jeko               1   Current   ST1            45   Standard
mefi               2   Current   ST1          1546   Root
saki               3   Expired   ST0           587   Standard
```

Related topics

- "Clients" on page 10

lscontainer

Displays a list of filesets (containers) and their attributes.

Syntax

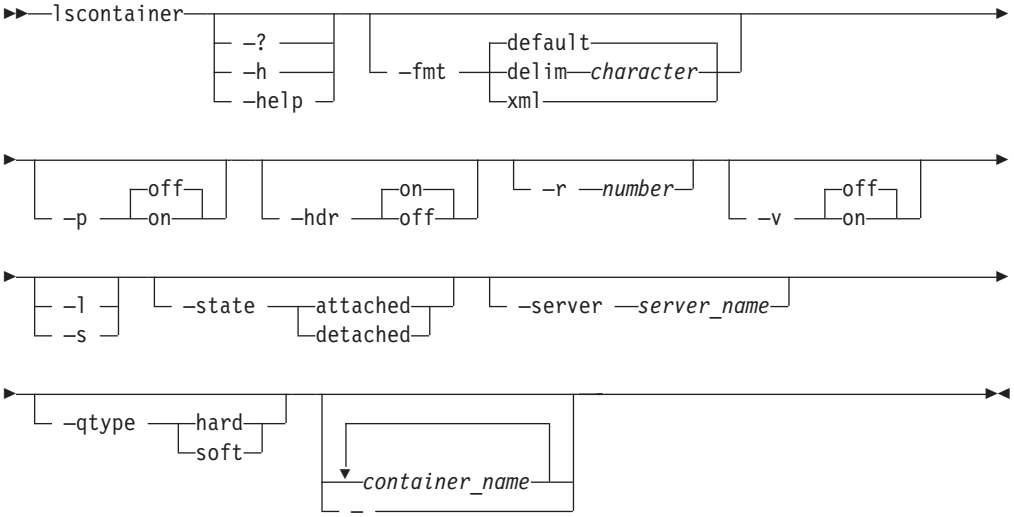

Parameters

–? | –h | –help

Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–fmt

Specifies the format of the output. You can specify one of the following values:

default

Specifies to display output in a tabular format using spaces as the delimiter between the columns. This is the default value. For example:

```
Name    Container State   Quota   Quota (MB)   Used (MB)   Used (%)
                          Type
=====================================================================
cnt_B   Attached         Soft            20          10          50

Threshold (%)   Most Recent Image

=====================================================================
           19   Feb 6, 2003 14:54:15 PM
```

delim *character*

Specifies to display output in a tabular format using the specified character to separate the columns. If you use a shell metacharacter (for example, * or \t) as the delimiting character, enclose the character in single quotation marks (') or double quotation marks ("). A blank space is not a valid character. For example:

```
cnt_B,Attached,Soft,20 MB,10 MB,50%,19%,Feb 6, 2003 14:54:15 PM
```

xml Specifies to display output using XML format, for example:

```
<IRETURNVALUE>
<INSTANCE CLASSNAME="STC_Container">
<PROPERTY NAME="Name" TYPE="string"><VALUE TYPE="string">cnt_B</VALUE>
</PROPERTY>
```

```
<PROPERTY NAME="State" TYPE="uint32"><VALUE TYPE="uint32">1</VALUE>
</PROPERTY>
<PROPERTY NAME="Quota" TYPE="uint64"><VALUE TYPE="uint64">20</VALUE>
</PROPERTY>
<PROPERTY NAME="IsHardQuota" TYPE="boolean"><VALUE TYPE="boolean">false
</VALUE></PROPERTY>
<PROPERTY NAME="AlertPercentage" TYPE="uint16"><VALUE TYPE="uint16">19
</VALUE></PROPERTY>
<PROPERTY NAME="SizeAllocated" TYPE="uint64"><VALUE TYPE="uint64">10
</VALUE></PROPERTY>
<PROPERTY NAME="SizeAllocatedPercentage" TYPE="uint16"><VALUE
TYPE="uint16">50</VALUE></PROPERTY>
<PROPERTY NAME="LastPITCopyDate" TYPE="datetime"><VALUE TYPE="datetime">
20030206145415&#46;000000&#43;000</VALUE></PROPERTY>
</INSTANCE>
</IRETURNVALUE>
```

–p Specifies whether to display one page of text at a time or all text at once.

off Displays all text at one time. This is the default value when the **tanktool** command is run in single-shot mode.

on Displays one page of text at time. Pressing any key displays the next page. This is the default value when the **tanktool** command is run in interactive mode.

–hdr
Specifies whether to display the table header.

on Displays the table header. This is the default value.

off Does not display the table header.

–r *number*
Specifies the number of rows per page to display when the **–p** parameter is on. The default is 24 rows. You can specify a value from 1 to 100.

–v Specifies whether to enable verbose mode.

off Disables verbose mode. This is the default value.

on Enables verbose mode.

–l Displays additional information about the filesets.

–s Displays only the names of the filesets.

–state
Lists only the filesets that match the specified state. Possible states are:

attached
Lists only the filesets that are attached to the global namespace.

detached
Lists only the filesets that are not attached to the global namespace.

–server *server_name*
Lists only filesets assigned to the specified Metadata server.

–qtype
Lists only the filesets (containers) that match the specified quota type. Possible types are:

hard Lists only the filesets that have hard quotas.

soft Lists only the filesets that have soft quotas.

container_name
Specifies the names of one or more filesets (or containers) to list. If not

specified, this command lists all the filesets. If not specified, this command displays a list of all filesets attached to the global namespace.

– Specifies that you want to read the names of one or more filesets to list from stdin (for example, – << /work/cnt_list.txt).

Description

Prerequisite: You must be logged in to the engine hosting the master Metadata server to run this command.

When you specify multiple parameters in this command, these behaviors apply:

- Flags work as AND coordinators.
- Objects work as OR coordinators.
- When flags and objects are combined, the they are listed in this manner: WHERE flag_1 AND flag_2 AND ... AND flag_n AND (object_1 OR object_2 OR ... OR object_m).
- When flags or objects are not required and not specified, this command searches all possible values unless a default overrides this rule.

For example, if you specify –state, –qtype, and two fileset names, the search would be performed in the following manner: WHERE state AND quota_type AND (fileset_name_1 OR fileset_name_2).

If you do not specify a listing parameter, this command displays the following default information for each fileset:

- Name of the fileset.
- State of the fileset (Detached, Idle, or Inuse).
- Quota type.
- Quota size, in MB.
- Amount of space used, in MB.
- Threshold, as a percentage of space used.
- Timestamp of the most recent FlashCopy image.

Note: If the quota is not set, the values for the amount of space used (Used %) and threshold (Threshold %) are displayed as 0 and have no meaning. Because the threshold value is based on the amount of space used, if the amount of space used is not defined, then any value specified for the threshold value will have no meaning and no alert will be sent if the displayed value is exceeded.

If you specify the **–l** listing parameter, this command displays the following information in addition to the default information for each fileset:

- Metadata server that hosts the fileset.
- State of the Metadata server.
- Number of FlashCopy images that exist for this fileset.
- Attach point.
- Directory name.
- Directory path.
- Parent fileset.
- Number of nested filesets (children).
- Fileset description.

Examples

Listing filesets The following example lists all the available information about all filesets in the global namespace.

```
tanktool> lscontainer -l

Name    Container State   Quota    Quota (MB)   Used (MB)   Used (%)
                          Type
=================================================================
ROOT    Attached          Soft             20          10          50
cnt_B   Detached          Hard             88         100          48
cnt_C   Attached          Hard            120         400          45

Threshold (%)   Most Recent Image

======================================
           19   -
           70   Nov 6, 2003 20:01:48 PM
            2   Mar 7, 2003 10:36:37 PM
```

Related topics

- "attachcontainer" on page 166
- "chcontainer" on page 174
- "detachcontainer" on page 183
- "mkcontainer" on page 230
- "rmcontainer" on page 251

lsdrfile

Displays a list of system-metadata disaster-recovery files.

Syntax

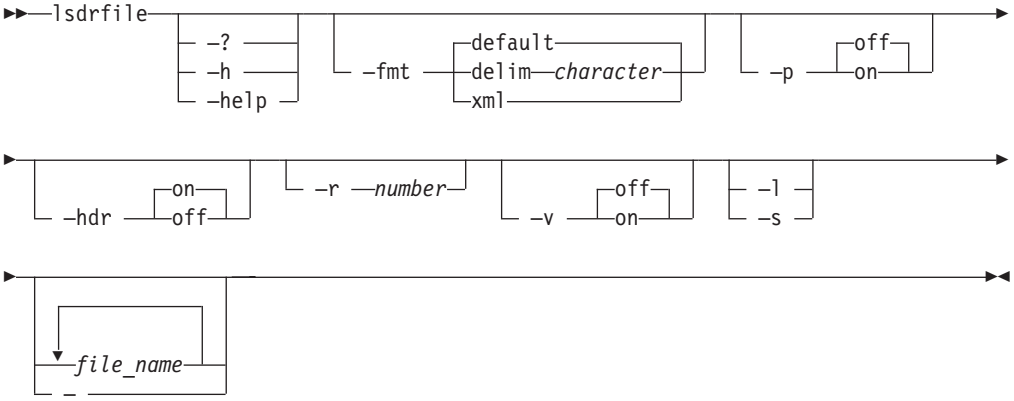

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–fmt
Specifies the format of the output. You can specify one of the following values:

default
Specifies to display output in a tabular format using spaces as the delimiter between the columns. This is the default value. For example:

```
Name          Date and Time            Size (KB)
================================================
SysMetadata   Jan 3, 2003 4:35:46 PM       10000
```

delim *character*
Specifies to display output in a tabular format using the specified character to separate the columns. If you use a shell metacharacter (for example, * or \t) as the delimiting character, enclose the character in single quotation marks (') or double quotation marks ("). A blank space is not a valid character. For example:

```
SysMetadata,Jan 3, 2003 4:35:46 PM,10000 KB
```

xml Specifies to display output using XML format, for example:

```
<IRETURNVALUE>
<INSTANCE CLASSNAME="STC_SystemMDRAid">
<PROPERTY NAME="Name" TYPE="string"><VALUE TYPE="string">SysMetadata
</VALUE></PROPERTY>
<PROPERTY NAME="InstallDate" TYPE="datetime"><VALUE TYPE="datetime">
20030103043546&#46;000000&45;420</VALUE></PROPERTY>
<PROPERTY NAME="Size" TYPE="uint64"><VALUE TYPE="uint64">10000</VALUE>
</PROPERTY>
h</INSTANCE>
</IRETURNVALUE>
```

–p Specifies whether to display one page of text at a time or all text at once.

off Displays all text at one time. This is the default value when the **tanktool** command is run in single-shot mode.

on Displays one page of text at time. Pressing any key displays the next page. This is the default value when the **tanktool** command is run in interactive mode.

–hdr
Specifies whether to display the table header.

on Displays the table header. This is the default value.

off Does not display the table header.

–r *number*
Specifies the number of rows per page to display when the **–p** parameter is on. The default is 24 rows. You can specify a value from 1 to 100.

–v Specifies whether to enable verbose mode.

off Disables verbose mode. This is the default value.

on Enables verbose mode.

–l Displays the default information for the system-metadata disaster-recovery files.

–s Displays only the names of the system-metadata disaster-recovery files.

file_name
Specifies the name of one or more system-metadata disaster-recovery files to display. If not specified, this command displays information for all system-metadata disaster-recovery files.

– Specifies the file from which this commands reads the names of one or more system-metadata disaster-recovery files to display.

Description

Note: This command operates only against the local engine.

If you do not specify a listing parameter or if you specify the **–l** listing parameter, this command displays the following default information for each system-metadata disaster-recovery file:
- Name of the disaster-recovery file.
- Date and time the disaster-recovery file was created.
- Size of the disaster-recovery file, in KB.

Examples

List all disaster-recovery files The following example displays information for all disaster-recovery files.

```
tanktool> lsdrfile

Name          Date and Time            Size (KB)
=============================================
SysMetadata   Jan 3, 2003 4:35:46 PM       10000
```

Related topics

- "builddrscript" on page 168
- "mkdrfile" on page 233
- "rmdrfile" on page 252

lsengine

Displays a list of storage engines and their attributes.

Syntax

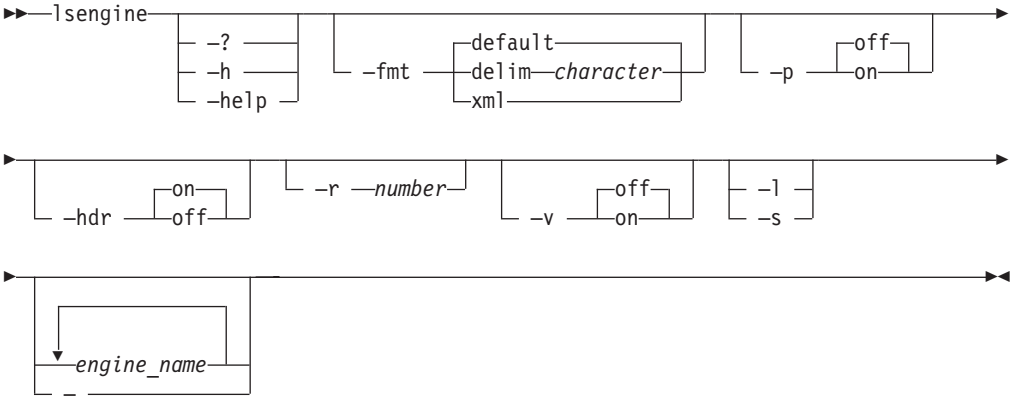

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–fmt
Specifies the format of the output. You can specify one of the following values:

default
Specifies to display output in a tabular format using spaces as the delimiter between the columns. This is the default value. For example:

```
Name   Engine IP      Boot State   Temp      Fans     Voltage
============================================================
ST0    129.42.16.99   In POST      Warning   Normal   Normal
```

delim *character*
Specifies to display output in a tabular format using the specified character to separate the columns. If you use a shell metacharacter (for example, * or \t) as the delimiting character, enclose the character in single quotation marks (') or double quotation marks ("). A blank space is not a valid character. For example:

```
ST0,129.42.16.99,In POST,Warning,Normal,Normal
```

xml Specifies to display output using XML format, for example:

```
<IRETURNVALUE>
<INSTANCE CLASSNAME="STC_ComputerSystem">
<PROPERTY NAME="CurrentState" TYPE="uint32"><VALUE TYPE="uint32">1
</VALUE></PROPERTY>
<PROPERTY NAME="Name" TYPE="string"><VALUE TYPE="string">ST0</VALUE>
</PROPERTY>
</INSTANCE>
</IRETURNVALUE>
```

–p Specifies whether to display one page of text at a time or all text at once.

off Displays all text at one time. This is the default value when the **tanktool** command is run in single-shot mode.

on Displays one page of text at time. Pressing any key displays the next page. This is the default value when the **tanktool** command is run in interactive mode.

–hdr
Specifies whether to display the table header.

on Displays the table header. This is the default value.

off Does not display the table header.

–r *number*
Specifies the number of rows per page to display when the **–p** parameter is on. The default is 24 rows. You can specify a value from 1 to 100.

–v Specifies whether to enable verbose mode.

off Disables verbose mode. This is the default value.

on Enables verbose mode.

–l Displays the default information for each engine.

–s Displays only the IP address and name of each engine.

engine_name
Specifies the names of one or more storage engines to display. If not specified, this command displays information for all storage engines.

– Specifies the file from which this command reads the names of one or more storage engines to display.

Description

If you specify the **–s** listing parameter, this command displays the IP address and name of each engine.

If you do not specify a listing parameter or if you specify the **–l** listing parameter, this command displays the following default information for each engine:

- Engine IP address.
- Engine name.
- Boot state. Possible values are:

Before POST
The engine is powered on but has not started the power-on self test (POST).

Booted Flash
The engine has read the System BIOS but has not started loading the operating system.

Booting OS
The engine has stated but not completed loading the operating system.

CPUs Held in Reset
The engine has been reset after a hardware fault.

In POST
The engine is running the POST.

In OS The engine is running in the normal state.

Stopped in POST (error detected)
The engine is powered on but has not completed the POST due to an error.

Unknown / Power Off
The boot state is unknown. The engine could be powered off.

- Temperature state. Possible values are:

  – The temperature thresholds are not set.

  Normal
  The temperatures of all components in all engine are below the Warning threshold.

  Warning
  The temperature of one or more engine components are above the Warning threshold.

  Error The temperature of one or more engine components is higher than the Soft Shutdown threshold.

  Unknown
  The RSA card could not be accessed.

- Fan state. Possible values are:

  Normal
  All fans are operating above 15% of its fan-speed capacity.

  Warning
  One or more fans are operating below 15% of its fan-speed capacity.

  Error The RSA card could not be accessed.

- Voltage state. Possible values are:

  – The voltage thresholds are not set.

  Normal
  The voltages of all components in all engines are above the Warning Low Voltage threshold and below the Warning High Voltage threshold.

  Warning
  The voltage of one or more engine components is below the Warning Low Voltage threshold or above the Warning High Voltage threshold.

  Error The RSA card could not be accessed.

Examples

List engines The following example displays information about two engines: ST0 and ST1.

```
tanktool> lsengine ST1 ST2

Name   Engine IP      Boot State   Temp      Fans     Voltage
============================================================
ST0    129.42.16.99   In POST      Warning   Normal   Normal
ST1    129.42.16.98   In OS        Normal    Normal   Normal
```

Related topics

- "startengine" on page 273
- "statengine" on page 282
- "stopengine" on page 291

lsimage

Displays a list of FlashCopy images in all filesets or a specified fileset.

Syntax

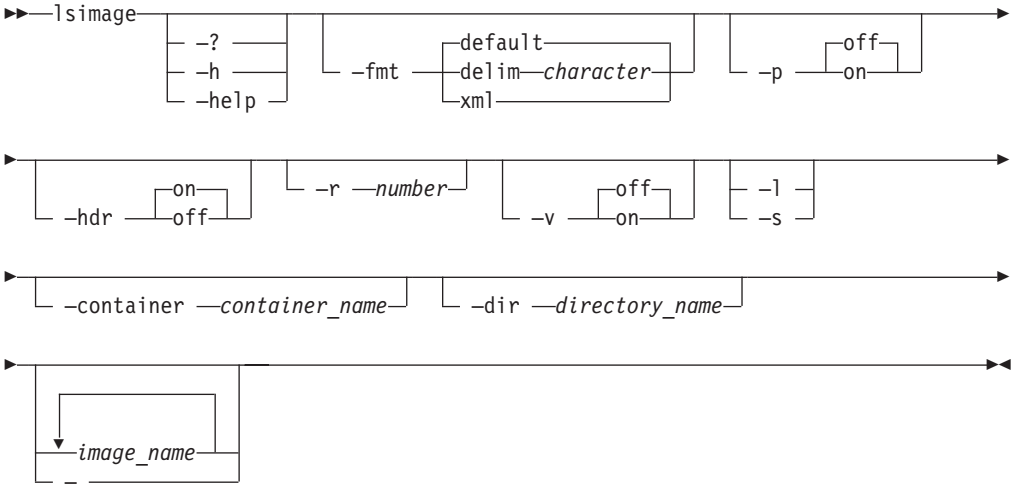

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–fmt
Specifies the format of the output. You can specify one of the following values:

default
Specifies to display output in a tabular format using spaces as the delimiter between the columns. This is the default value. For example:

```
Name        Container   Directory Name   Date
============================================================
FCA_image   cnt_A       personnel        Dec 3, 2003 5:48:46 PM
```

delim *character*
Specifies to display output in a tabular format using the specified character to separate the columns. If you use a shell metacharacter (for example, * or \t) as the delimiting character, enclose the character in single quotation marks (') or double quotation marks ("). A blank space is not a valid character. For example:

```
FCA_image,cnt_A,personnel,Dec 3, 2003 5:48:46 PM
```

xml Specifies to display output using XML format, for example:

```
<IRETURNVALUE>
<INSTANCE CLASSNAME="STC_PitImage">
<PROPERTY NAME="ContainerName" TYPE="string"><VALUE TYPE="string">cnt_A<
/VALUE></PROPERTY>
<PROPERTY NAME="Name" TYPE="string"><VALUE TYPE="string">FCA_imag<
/VALUE></PROPERTY>
<PROPERTY NAME="InstallDate" TYPE="datetime"><VALUE TYPE="datetime">
20031203174846&#46;000009&#45;420</VALUE></PROPERTY>
<PROPERTY NAME="DirectoryName" TYPE="string"><VALUE TYPE="string">
personnel</VALUE></PROPERTY>
</INSTANCE>
</IRETURNVALUE>
```

–p Specifies whether to display one page of text at a time or all text at once.

off Displays all text at one time. This is the default value when the **tanktool** command is run in single-shot mode.

on Displays one page of text at time. Pressing any key displays the next page. This is the default value when the **tanktool** command is run in interactive mode.

–hdr
Specifies whether to display the table header.

on Displays the table header. This is the default value.

off Does not display the table header.

–r *number*
Specifies the number of rows per page to display when the **–p** parameter is on. The default is 24 rows. You can specify a value from 1 to 100.

–v Specifies whether to enable verbose mode.

off Disables verbose mode. This is the default value.

on Enables verbose mode.

–l Displays additional information about the FlashCopy images.

–s Displays only the name of the FlashCopy images and the fileset associated with each.

–container *container_name*
Specifies the name of the fileset for which to list all FlashCopy images (or containers). If not specified, this command lists FlashCopy images in all filesets.

–dir *directory_name*
Specifies the directory under the .flashcopy directory from which to list the FlashCopy image. This directory can be up to 256 characters in length and must not contain backslash (\) or forward slash (/) characters as delimiters. If not specified, this command lists the FlashCopy images in all directories under the .flashcopy directory.

image_name
Specifies the names of one or more FlashCopy images to display. If not specified, this command lists all FlashCopy images in the specified fileset and directory.

– Specifies that you want to read the names of one or more FlashCopy images to display from stdin (for example, – << /work/image_list.txt).

Description

Prerequisite: You must be logged in to the engine hosting the master Metadata server to run this command.

When you specify multiple parameters in this command, these behaviors apply:

- Flags work as AND coordinators.
- Objects work as OR coordinators.
- When flags and objects are combined, the they are listed in this manner: WHERE flag_1 AND flag_2 AND ... AND flag_n AND (object_1 OR object_2 OR ... OR object_m).

- When flags or objects are not required and not specified, this command searches all possible values unless a default overrides this rule.

For example, if you specify –container, –dir, and two FlashCopy image names, the search would be performed in the following manner: WHERE fileset AND directory AND (image_name_1 OR image_name_2).

The most-recently created FlashCopy image are listed first.

If no parameters are specified, this command lists all FlashCopy images in all filesets and directories.

If you do not specify a listing parameter, this command displays the following default information for each FlashCopy image:
- FlashCopy image name.
- Fileset name.
- Directory name.
- Timestamp when the FlashCopy image was created.

If you specify the **–l** listing parameter, this command displays the following information in addition to the default information for each FlashCopy image:
- Description.

Examples

Listing FlashCopy images The following example lists all the available information about the FlashCopy images.

```
tanktool> lsimage –l

Name        Container   Directory Name   Date                    Description
================================================================================
FCC_image   cnt_C       payroll          Dec 3, 2003 5:59:14 PM   Payroll data
FCB_image   cnt_B       website          Dec 3, 2003 5:53:26 PM   Web site data
FCA_image   cnt_A       personnel        Dec 3, 2003 5:48:46 PM   Personnel data
```

Related topics

- “mkimage” on page 234
- “reverttoimage” on page 249
- “rmimage” on page 253

lslun

Lists the LUNs that are accessible from SAN File System.

Syntax

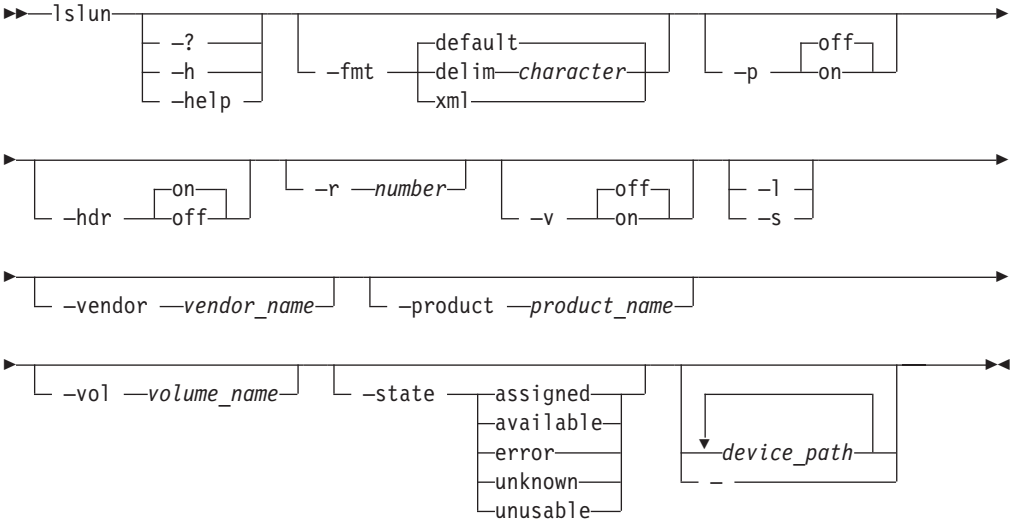

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–fmt
Specifies the format of the output. You can specify one of the following values:

default
Specifies to display output in a tabular format using spaces as the delimiter between the columns. This is the default value. For example:

```
OS Device Path   LUN ID   Engine WWN                Vendor  Product
=========================================================================
/dev/sda              2   12:34:56:78:91:23:45:67   IBM     2105F20

Size (MB)   Volume   State
============================
    47000   Vol1     Assigned
```

delim *character*
Specifies to display output in a tabular format using the specified character to separate the columns. If you use a shell metacharacter (for example, * or \t) as the delimiting character, enclose the character in single quotation marks (') or double quotation marks ("). A blank space is not a valid character. For example:

```
/dev/sda,2,12:34:56:78:91:23:45:67,IBM,2105F20,47000,Vol1,Assigned
```

xml Specifies to display output using XML format, for example:

```
<IRETURNVALUE>
<INSTANCE CLASSNAME="STC_AvailableLUNs">
<PROPERTY NAME="LUNID" TYPE="uint64"><VALUE TYPE="uint64">2</VALUE>
</PROPERTY>
<PROPERTY NAME="NodeWWN" TYPE="string"><VALUE TYPE="string">
12&58;34&58;56&58;78&58;91&58;23&58;45&58;67</VALUE></PROPERTY>
```

```
<PROPERTY NAME="Vendor" TYPE="string"><VALUE TYPE="string">IBM</VALUE>
</PROPERTY>
<PROPERTY NAME="Product" TYPE="string"><VALUE TYPE="string">2105F20
</VALUE></PROPERTY>
<PROPERTY NAME="Size" TYPE="uint64"><VALUE TYPE="uint64">47000</VALUE>
</PROPERTY>
<PROPERTY NAME="State" TYPE="uint64"><VALUE TYPE="uint64">2</VALUE>
</PROPERTY>
<PROPERTY NAME="VolumeName" TYPE="string"><VALUE TYPE="string"></VALUE>
</PROPERTY>
<PROPERTY NAME="DeviceID" TYPE="string"><VALUE TYPE="string">
&47;dev&47;sda</VALUE></PROPERTY>
</INSTANCE>
</IRETURNVALUE>
```

–p Specifies whether to display one page of text at a time or all text at once.

off Displays all text at one time. This is the default value when the **tanktool** command is run in single-shot mode.

on Displays one page of text at time. Pressing any key displays the next page. This is the default value when the **tanktool** command is run in interactive mode.

–hdr
Specifies whether to display the table header.

on Displays the table header. This is the default value.

off Does not display the table header.

–r *number*
Specifies the number of rows per page to display when the **–p** parameter is on. The default is 24 rows. You can specify a value from 1 to 100.

–v Specifies whether to enable verbose mode.

off Disables verbose mode. This is the default value.

on Enables verbose mode.

–l Displays default information about the LUNs.

–s Displays default information about the LUNs.

–vendor *vendor_name*
Displays information for only the LUNs that match the specified vendor name.

–product *product_name*
Displays information for only the LUNs that match the specified product name.

–vol *volume_name*
Displays information for only the LUNs that match the specified SAN File System volume name.

–state
Displays information for only the LUNs that match the specified state. Possible states are:

assigned
Displays only those LUNs that are assigned to a storage pool and usable.

available
Displays only those LUNs that are visible to SAN File System and usable but are not assigned to a storage pool.

error Displays only those LUNs for which an error occurred while determining their availability.

unknown
Displays only those LUNs whose availability could not be determined because the Metadata server is not running.

unusable
Displays only those LUNs that are unsuitable to be a volume. A LUN is considered unsuitable if it is read-only, if its availability from all Metadata servers in the cluster is inconsistent, or if the LUN does not have an associated raw device.

device_path
Specifies one or more operating-system device paths for which to display LUNs.

– Specifies that you want to read one or more operating-system device paths for which to display LUNs from stdin (for example, – << /work/device_paths_list.txt).

Description

Note: This command operates only against the local engine.

If you do not specify any filtering parameters, this command lists all the LUNs that are accessible to SAN File System, including LUNs that have not been added to SAN File System, SAN File System volumes, and unusable LUNs.

This command displays the following default information for each LUN:
- Operating-system device path.
- LUN ID.
- Worldwide name of the engine.
- Vendor.
- Product.
- Size of the device, in MB.
- Volume name, if the LUN is assigned in SAN File System.
- State of the LUN.

Examples

Display LUNs The following example lists all the attributes of all the available LUNs:

```
tanktool> lslun

OS Device Path   LUN ID   Storage Engine WWN       Vendor  Product
=====================================================================
/dev/sda             2   12:34:56:78:91:23:45:67   IBM     2105F20
/dev/sda7            3   89:12:34:56:78:91:23:45   IBM     2105F20
/dev/sda8            4   88:19:34:56:78:91:23:41   IBM     2105800

Size (MB)   Volume   State
==============================
    47000   Vol1     Assigned
       12   -        Available
    80000   -        Unusable
```

Related topics

- "lsvol" on page 227

lspolicy

Lists the active and inactive policies.

Syntax

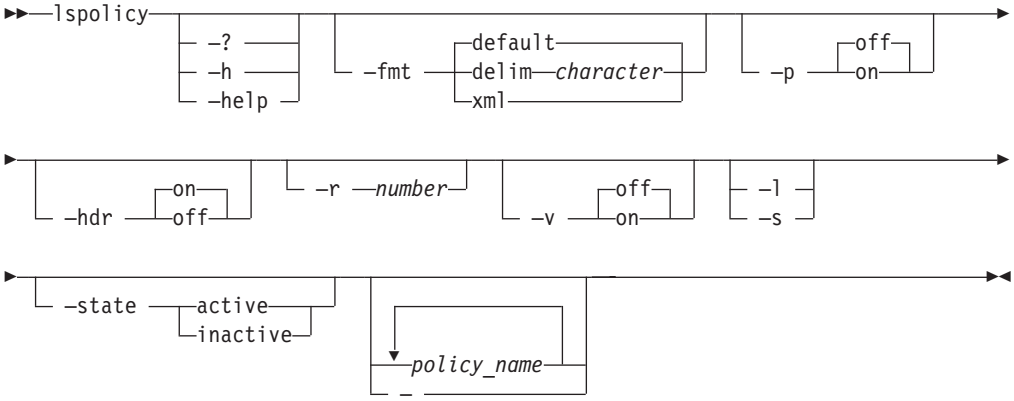

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–fmt
Specifies the format of the output. You can specify one of the following values:

default
Specifies to display output in a tabular format using spaces as the delimiter between the columns. This is the default value. For example:

```
Name         State   Last Active                Modified
==========================================================================
policy1      Active  Jul 19, 2003 14:00:34 PM   Jul 19, 2003 13:59:34 PM

Description
================
Active policy
```

delim *character*
Specifies to display output in a tabular format using the specified character to separate the columns. If you use a shell metacharacter (for example, * or \t) as the delimiting character, enclose the character in single quotation marks (') or double quotation marks ("). A blank space is not a valid character. For example:

```
policy1,Active,Jul 19, 2003 14:00:34 PM,Jul 19, 2003 13:59:34 PM,Active
policy
```

xml Specifies to display output using XML format, for example:

```
<IRETURNVALUE>
<INSTANCE CLASSNAME="STC_PolicySet">
<PROPERTY NAME="Name" TYPE="string"><VALUE TYPE="string">policy1</VALUE>
</PROPERTY>
<PROPERTY NAME="State" TYPE="uint16"><VALUE TYPE="uint16">1</VALUE>
</PROPERTY>
<PROPERTY NAME="Description" TYPE="string"><VALUE TYPE="string">
Active&#32;policy</VALUE></PROPERTY>
<PROPERTY NAME="LastModificationDate" TYPE="datetime"><VALUE
TYPE="datetime">20030719135934&#45;000000&#45;420</VALUE></PROPERTY>
```

```
<PROPERTY NAME="LastActiveDate" TYPE="datetime"><VALUE TYPE="datetime">
20030719140034&#45;000000&#45;420</VALUE></PROPERTY>
</INSTANCE>
</IRETURNVALUE>
```

–p Specifies whether to display one page of text at a time or all text at once.

off Displays all text at one time. This is the default value when the **tanktool** command is run in single-shot mode.

on Displays one page of text at time. Pressing any key displays the next page. This is the default value when the **tanktool** command is run in interactive mode.

–hdr
Specifies whether to display the table header.

on Displays the table header. This is the default value.

off Does not display the table header.

–r *number*
Specifies the number of rows per page to display when the **–p** parameter is on. The default is 24 rows. You can specify a value from 1 to 100.

–v Specifies whether to enable verbose mode.

off Disables verbose mode. This is the default value.

on Enables verbose mode.

–l Displays the default information about each policy.

–s Displays only the name of each policy.

–state active | inactive
Specifies the state of the policies to display. You can specify one of the following values:

active Displays only the active policy.

inactive
Displays only the inactive policies.

policy_name
Specifies the names of one or more policies to list.

– Specifies that you want to read the names of one or more policies to list from stdin (for example, – << /work/policies_list.txt).

Description

Prerequisite: You must be logged in to the engine hosting the master Metadata server to run this command.

When you specify multiple parameters in this command, these behaviors apply:

- Flags work as AND coordinators.
- Objects work as OR coordinators.
- When flags and objects are combined, the they are listed in this manner: WHERE flag_1 AND flag_2 AND ... AND flag_n AND (object_1 OR object_2 OR ... OR object_m).
- When flags or objects are not required and not specified, this command searches all possible values unless a default overrides this rule.

For example, if you specify –state and two policy names, the search would be performed in the following manner: WHERE state AND (policy_name_1 OR policy_name_2).

If you do not specify a listing parameter or if you specify the **–l** listing parameter, this command displays the following information for each policy:

- Name of the policy.
- State of the policy.
- Timestamp when the policy was last active, if applicable, based on coordinated universal time (UTC).
- Timestamp when the policy was last modified, based on coordinated universal time (UTC).
- Policy description.

Examples

Display policies The following example displays all the available information about all the policies:

```
tanktool> lspolicy –l

Name         State      Last Active                Modified
============================================================================
policy1      Active     Jul 19, 2003 14:00:34 PM   Jul 19, 2003 13:59:34 PM
testpolicy   Inactive   Mar 18, 2003 10:30:17 PM   Mar 18, 2003 10:06:25 PM
DEFAULT      Inactive   Jan 11, 2003 16:36:24 PM   Jan 11, 2003 15:36:14 PM

Description
===================
Active policy
Test policy
Default Policy
```

Display the rules of the active policy The following example displays the rules of the active policy:

```
$tanktool lspolicy –s –state active –hdr off | tanktool catpolicy –
VERSION 1
rule 'stgRule1' set stgpool 'pool1' for fileset ('cnt_A')
rule 'stgRule2' set stgpool 'pool2' where NAME like '%.doc'
rule 'stgRule3' set stgpool 'pool3' where DAYOFWEEK(CREATION_DATE) == 1
rule 'stgRule4' set stgpool 'pool4' where USER_ID <= 100
```

Related topics

- “catpolicy” on page 171
- “mkpolicy” on page 236
- “rmpolicy” on page 255
- “usepolicy” on page 296

lspool

Displays a list of the existing storage pools and their attributes.

Syntax

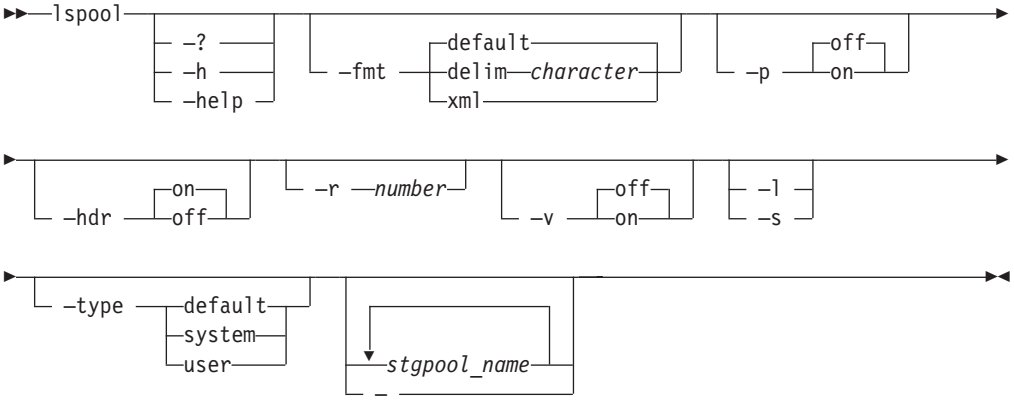

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–fmt
Specifies the format of the output. You can specify one of the following values:

default
Specifies to display output in a tabular format using spaces as the delimiter between the columns. This is the default value. For example:

```
Name    Type   Size (MB)  Used (MB)   Used (%)  Threshold (%)   Volumes
=========================================================================
Pool1   User        1400        575         41             90         2
```

delim *character*
Specifies to display output in a tabular format using the specified character to separate the columns. If you use a shell metacharacter (for example, * or \t) as the delimiting character, enclose the character in single quotation mark (') or double quotation mark ("). A blank space is not a valid character. For example:

```
Pool1,User,
```

xml Specifies to display output using XML format, for example:

```
<IRETURNVALUE>
<INSTANCE CLASSNAME="STC_StoragePool">
<PROPERTY NAME="Name" TYPE="string"><VALUE TYPE="string">Pool1</VALUE>
</PROPERTY>
<PROPERTY NAME="PoolType" TYPE="uint32"><VALUE TYPE="uint32">0</VALUE>
</PROPERTY>
<PROPERTY NAME="AlertPercentage" TYPE="uint16"><VALUE TYPE="uint16">90
</VALUE></PROPERTY>
<PROPERTY NAME="Size" TYPE="uint64"><VALUE TYPE="uint64">1400</VALUE>
</PROPERTY>
<PROPERTY NAME="SizeAllocated" TYPE="uint64"><VALUE TYPE="uint64">575
</VALUE></PROPERTY>
<PROPERTY NAME="SizeAllocatedPercentage" TYPE="uint16"><VALUE
TYPE="uint16">41</VALUE></PROPERTY>
```

```
<PROPERTY NAME="NumberOfVolumes" TYPE="uint32"><VALUE TYPE="uint32">2
</VALUE></PROPERTY>
</INSTANCE>
</IRETURNVALUE>
```

–p Specifies whether to display one page of text at a time or all text at once.

off Displays all text at one time. This is the default value when the **tanktool** command is run in single-shot mode.

on Displays one page of text at time. Pressing any key displays the next page. This is the default value when the **tanktool** command is run in interactive mode.

–hdr
Specifies whether to display the table header.

on Displays the table header. This is the default value.

off Does not display the table header.

–r *number*
Specifies the number of rows per page to display when the **–p** parameter is on. The default is 24 rows. You can specify a value from 1 to 100.

–v Specifies whether to enable verbose mode.

off Disables verbose mode. This is the default value.

on Enables verbose mode.

–l Displays additional information about the storage pools.

–s Displays only the names of the storage pools.

–type
Displays information for only those storage pools that match the specified type. If not specified, all storage pools are displayed. Possible types are:

default
Displays information for the DEFAULT storage pool.

system
Displays information for the SYSTEM storage pool.

user Displays information for all user storage pools.

stgpool_name
Specifies the names of one or more storage pools to display.

– Specifies that you want to read the names of one or more storage pools to display from stdin (for example, – << /work/stgpools_list.txt).

Description

Prerequisite: You must be logged in to the engine hosting the master Metadata server to run this command.

If you do not specify a listing parameter, this command displays the following default information for each storage pool:
- Name of the storage pool.
- Type of storage pool.
- Size of the storage pool, in MB.
- Amount of space used, in MB.
- Percent of space used.
- Alert threshold.

- Number of volumes in the storage pool.

If you specify the **–l** listing parameter, this command displays the following information in addition to the default information for each storage pool:
- Partition size, in MB (16, 64, or 256).
- Allocation size, in KB (auto, 4 or 128).
- Pool description.

Examples

Display a long listing The following example lists all the attributes of all the storage pools:

```
tanktool> lspool –l

Name      Type     Size (MB)  Used (MB)   Used (%)  Threshold (%)   Volumes
=============================================================================
Pool1     User          1400        575         41             90         2
Pool2     User          2000       1000         50             90         3
DEFAULT   Default      10000       2500         25             80        10
SYSTEM    System       10000       7500         75             80        10

Partition Size (MB)   Allocation Size (KB)   Description
=========================================================
                 16                      2   Accounting
                 16                      3   Engineering
                 64                    128   DEFAULT pool
                 64                    128   SYSTEM pool
```

Related topics

- “chpool” on page 176
- “mkpool” on page 238
- “rmpool” on page 256
- “setdefaultpool” on page 262

lsproc

Displays a list of long-running processes that are not yet complete and their attributes.

Syntax

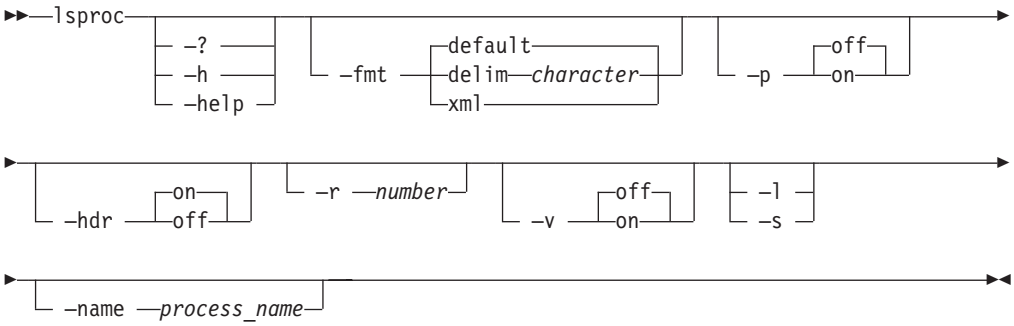

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–fmt
Specifies the format of the output. You can specify one of the following values:

default
Specifies to display output in a tabular format using spaces as the delimiter between the columns. This is the default value. For example:

```
ID   Name            Started
=============================================
5    CHECKMETADATA   Feb 11, 2003 9:29:56:13 PM
```

delim *character*
Specifies to display output in a tabular format using the specified character to separate the columns. If you use a shell metacharacter (for example, * or \t) as the delimiting character, enclose the character in single quotation mark (') or double quotation mark ("). A blank space is not a valid character. For example:

```
5,CHECKMETADATA,Feb 11, 2003 9:29:56:13 PM
```

xml Specifies to display output using XML format, for example:

```
<IRETURNVALUE>
<INSTANCE CLASSNAME="STC_AdminProcesses">
<PROPERTY NAME="Id" TYPE="uint64"><VALUE TYPE="uint64">5</VALUE>
</PROPERTY>
<PROPERTY NAME="Command" TYPE="string"><VALUE TYPE="uint32">STOPSERVER
</VALUE></PROPERTY>
<PROPERTY NAME="InstallDate" TYPE="datetime"><VALUE TYPE="datetime">
20030211092956&#46;000000&#45;420</VALUE></PROPERTY>
</INSTANCE>
</IRETURNVALUE>
```

–p Specifies whether to display one page of text at a time or all text at once.

off Displays all text at one time. This is the default value when the **tanktool** command is run in single-shot mode.

on Displays one page of text at time. Pressing any key displays the next page. This is the default value when the **tanktool** command is run in interactive mode.

–hdr
Specifies whether to display the table header.

on Displays the table header. This is the default value.

off Does not display the table header.

–r *number*
Specifies the number of rows per page to display when the **–p** parameter is on. The default is 24 rows. You can specify a value from 1 to 100.

–v Specifies whether to enable verbose mode.

off Disables verbose mode. This is the default value.

on Enables verbose mode.

–l Displays additional information about the processes.

–s Displays only the identifier of each process.

–name *process_name*
Specifies the name of a process to list.

Description

Prerequisite: You must be logged in to the engine hosting the master Metadata server to run this command.

If you do not specify a listing parameter or if you specify the **–l** listing parameter, this command displays the following information for each process:
- Identifier of the process.
- Name of the process.
- Time that the process was started.

Examples

Lists long-running processes The following example lists the long-running processes currently in progress.

```
tanktool> lsproc

ID   Name            Started
=============================================
5    CHECKMETADATA   Feb 11, 2003 9:29:56:13 PM
7    STOPCLUSTER     Feb 10, 2003 9:30:20:24 PM
9    STARTSERVER     Feb 10, 2003 9:12:22:24 PM
```

Related topics

- "lsproc" on page 219

lsserver

Displays a lists of all Metadata servers in the cluster and their attributes if issued from the master Metadata server, or displays attributes about the local Metadata server if issued from a subordinate Metadata server.

Syntax

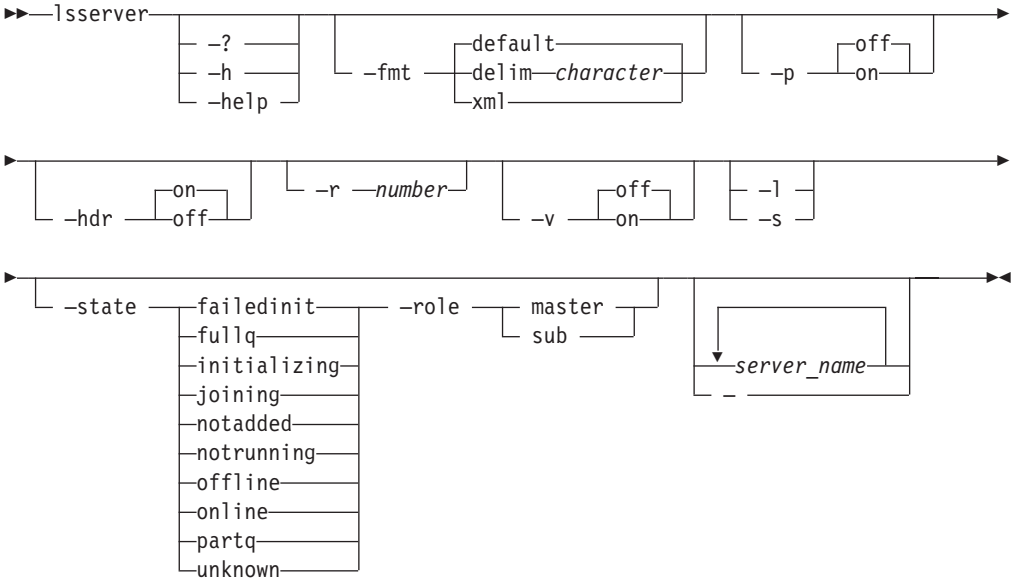

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–fmt
Specifies the format of the output. You can specify one of the following values:

default
Specifies to display output in a tabular format using spaces as the delimiter between the columns. This is the default value. For example:

```
Name   State    Relationship   Containers   Last Boot
=====================================================================
ST1    ONLINE   Subordinate    3            Dec. 31, 2002 5:00:34 PM
```

delim *character*
Specifies to display output in a tabular format using the specified character to separate the columns. If you use a shell metacharacter (for example, * or \t) as the delimiting character, enclose the character in single quotation mark (') or double quotation mark ("). A blank space is not a valid character. For example:

```
ST1;ONLINE;Subordinate;3;Dec. 31, 2002 5:00:34 PM
```

xml Specifies to display output using XML format, for example:

```
<IRETURNVALUE>
<INSTANCE CLASSNAME="STC_TankService">
<PROPERTY NAME="CurrentState" TYPE="uint32"><VALUE TYPE="uint32">1
</VALUE></PROPERTY>
<PROPERTY NAME="LastBootUpTime" TYPE="datetime"><VALUE TYPE="datetime">
20031231050034&#46;000000&#45;420</VALUE></PROPERTY>
```

```
<PROPERTY NAME="SystemName" TYPE="string"><VALUE TYPE="string">
DEFAULT_POOL</VALUE>ST1</PROPERTY>
<PROPERTY NAME="IsMaster" TYPE="boolean"><VALUE TYPE="boolean">false
</VALUE></PROPERTY>
<PROPERTY NAME="NumberOfContainers" TYPE="uint32"><VALUE TYPE="uint32">3
</VALUE></PROPERTY>
</INSTANCE>
</IRETURNVALUE>
```

–p Specifies whether to display one page of text at a time or all text at once.

off Displays all text at one time. This is the default value when the **tanktool** command is run in single-shot mode.

on Displays one page of text at time. Pressing any key displays the next page. This is the default value when the **tanktool** command is run in interactive mode.

–hdr
Specifies whether to display the table header.

on Displays the table header. This is the default value.

off Does not display the table header.

–r *number*
Specifies the number of rows per page to display when the **–p** parameter is on. The default is 24 rows. You can specify a value from 1 to 100.

–v Specifies whether to enable verbose mode.

off Disables verbose mode. This is the default value.

on Enables verbose mode.

–l Displays additional information about the Metadata servers.

–s Displays only the name of each Metadata server.

–state
Lists information for only those Metadata servers that match the specified state:

failedinit
Displays information for Metadata servers with the state FAILED INITIALIZATION. A fatal error occurred during initialization and the Metadata server is suspended. The Metadata server remains suspended so you can fix the problem that occurred during initialization.

fullq Displays information for Metadata servers with the state FULLY QUIESCENT. The Metadata server suspends all current client sessions and prevents new client sessions from being started. All file-data and metadata activity is suspended. The Metadata server flushes all dirty cached pages and commits them to disk. This state is used to produce self-consistent LUN-based backups from a known point-in-time.

initializing
Displays information for Metadata servers with the state INITIALIZING. The Metadata server is running, but has not yet opened its communication ports.

joining
Displays information for Metadata servers with the state JOINING. The Metadata server is joining the cluster.

notadded
Displays information for Metadata servers with the state NOT ADDED. The Metadata server is not part of the cluster. Because the master Metadata server does not know about Metadata servers that are not part of the cluster, this state is only available by logging into that Metadata server.

notrunning
Displays information for Metadata servers with the state NOT RUNNING. The Metadata server is part of the cluster but is not running and cannot perform any functions.

offline
Displays information for Metadata servers with the state OFFLINE. The Metadata server terminates all current client sessions and prevents new client sessions from being started. This state allows all Metadata server I/O to continue. This state restricts the Metadata server from client access.

online Displays information for Metadata servers with the state ONLINE. The Metadata server has fully initialized all of its subsystems, is a member of the cluster, and is serving client requests.

partq Displays information for Metadata servers with the state PARTLY QUIESCENT. The Metadata server prevents new client sessions from being started but allows file I/O to continue for existing clients. The Metadata server suspends clients making file metadata updates (such as changing permissions or creating new files) and system metadata activity. The Metadata server also flushes any dirty metadata pages in the client and Metadata server caches. This state allows you to produce dirty LUN-based backups without requiring clients and applications to stop. You can back up applications using direct I/O (such as database servers) (no caching) in this state.

unknown
Displays information for Metadata servers with the state UNKNOWN.

–role master | sub
Displays only information for those Metadata servers that match the specified role:

master
Displays information for the master Metadata server.

sub Displays information for subordinate Metadata servers.

server_name
Specifies the name of one or more names of Metadata servers to display. If not specified and if this command is issued from the master issued from the master Metadata server, lists information about all Metadata servers in the cluster.

– Specifies that you want to read the names of one or more of the Metadata servers to display from stdin (for example, – << /work/server_list.txt).

Description

Note:

If you run this command from an engine hosting a subordinate Metadata server, you can display information about only the local Metadata server. If

you run this command from the engine hosting the master Metadata server, you can display information about any Metadata server.

When you specify multiple parameters in this command, these behaviors apply:

- Flags work as AND coordinators.
- Objects work as OR coordinators.
- When flags and objects are combined, the they are listed in this manner: WHERE flag_1 AND flag_2 AND ... AND flag_n AND (object_1 OR object_2 OR ... OR object_m).
- When flags or objects are not required and not specified, this command searches all possible values unless a default overrides this rule.

For example, if you specify –state, –role, and two Metadata server names, the search would be performed in the following manner: WHERE state AND role AND (server_name_1 OR server_name_2).

If you do not specify a listing parameter, this command displays the following default information for each Metadata server:

- Name of the Metadata server.
- State of the Metadata server.
- Role of the Metadata server (Master or Subordinate).
- Number of filesets that are hosted by the Metadata server.
- Timestamp of the last boot.

If you specify the **–l** listing parameter, this command displays the following information in addition to the default information for each Metadata server:

- Timestamp of the last status change.
- State to which the Metadata server changed.
- Timestamp of the last target state change.
- Current time.
- Pending software version.

Examples

Display online Metadata servers The following example lists all the attributes of all the Metadata servers with the state ONLINE:

```
tanktool> lsserver –state online

Name   State    Relationship   Containers   Last Boot
=================================================================
ST1    ONLINE   Subordinate    3            Dec. 31, 2002 5:00:34 PM
ST2    ONLINE   Master         2            Dec. 31, 2002 5:00:24 PM
```

Related topics

- “Metadata server” on page 8
- “statserver” on page 286

lssnmpmgr

Displays a list of SNMP managers and their attributes.

Syntax

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

Description

Prerequisite: You must be logged in to the engine hosting the master Metadata server to run this command.

This command displays the following information for each SNMP manager:
- SNMP manager identifier.
- IP address.
- Port number.
- Version.
- Community.

Examples

Lists the SNMP managers The following example lists the SNMP managers.

```
tanktool> lssnmpmgr

ID     IP            Port   Version   Community
=============================================
 1     192.168.0.1   8192   V1        Public
 2     192.168.0.2   162    V2C       SNMPMgr2
```

Related topics

- "addsnmpmgr" on page 164
- "catlog" on page 169
- "clearlog" on page 180
- "lstrapsetting" on page 226
- "rmsnmpmgr" on page 257
- "settrap" on page 268
- "SNMP" on page 18

lstrapsetting

Displays a list of event types that currently generate an SNMP trap.

Syntax

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

Description

Prerequisite: You must be logged in to the engine hosting the master Metadata server to run this command.

This command displays the whether the following event types are enabled:
- Information
- Error
- Warning
- Severe

Examples

Sets the SNMP traps The following example sets SNMP traps for severe and warning events.

```
tanktool> lstrapsetting
Information       Disabled
Error             Enabled
Warning           Enabled
Severe            Enabled
```

Related topics

- “catlog” on page 169
- “clearlog” on page 180
- “lssnmpmgr” on page 225
- “settrap” on page 268

lsvol

Displays a list of available volumes.

Syntax

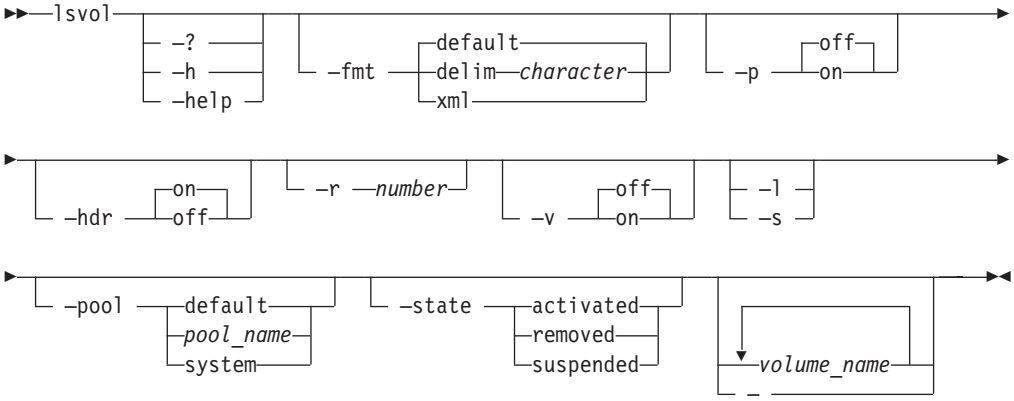

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–fmt
Specifies the format of the output. You can specify one of the following values:

default
Specifies to display output in a tabular format using spaces as the delimiter between the columns. This is the default value. For example:

```
Name   State           Pool        Size (MB)   Used (MB)   Used (%)
=========================================================================
Vol1   Activated       P_Accting         700         350         50
```

delim *character*
Specifies to display output in a tabular format using the specified character to separate the columns. If you use a shell metacharacter (for example, * or \t) as the delimiting character, enclose the character in single quotation marks (') or double quotation marks ("). A blank space is not a valid character. For example:

```
Vol1,Activated,P_Accting,700,350,50
```

xml Specifies to display output using XML format, for example:

```
<IRETURNVALUE>
<INSTANCE CLASSNAME="STC_Volume">
<PROPERTY NAME="StoragePoolName" TYPE="string"><VALUE TYPE="string">
P_Accting</VALUE></PROPERTY>
<PROPERTY NAME="Name" TYPE="string"><VALUE TYPE="string">Vol1</VALUE>
</PROPERTY>
<PROPERTY NAME="State" TYPE="uint32"><VALUE TYPE="uint32">0</VALUE>
</PROPERTY>
<PROPERTY NAME="Size" TYPE="uint64"><VALUE TYPE="uint64">700</VALUE>
</PROPERTY>
<PROPERTY NAME="SizeAllocated" TYPE="uint64"><VALUE TYPE="uint64">350
</VALUE></PROPERTY>
```

```
<PROPERTY NAME="SizeAllocatedPercentage" TYPE="uint16"><VALUE
TYPE="uint16">50</VALUE></PROPERTY>
</INSTANCE>
</IRETURNVALUE>
```

–p Specifies whether to display one page of text at a time or all text at once.

off Displays all text at one time. This is the default value when the **tanktool** command is run in single-shot mode.

on Displays one page of text at time. Pressing any key displays the next page. This is the default value when the **tanktool** command is run in interactive mode.

–hdr
Specifies whether to display the table header.

on Displays the table header. This is the default value.

off Does not display the table header.

–r *number*
Specifies the number of rows per page to display when the **–p** parameter is on. The default is 24 rows. You can specify a value from 1 to 100.

–v Specifies whether to enable verbose mode.

off Disables verbose mode. This is the default value.

on Enables verbose mode.

–l Displays additional information about the volumes.

–s Displays only the name of the volumes.

–pool *pool_name*
Specifies the name of the storage pool to display. You can specify one of the following values:

- **default** (for the default storage pool)
- User-defined storage pool name
- **system** (for the system storage pool)

–state
Displays only information for those storage pools that match the specified type. Possible types are:

activated
Displays information for all volumes with a status of Activated.

removed
Displays information for all volumes with a status of Being Removed.

suspended
Displays information for all volumes with a status of Suspended.

volume_name
Specifies the names of one or more volumes to list. If unspecified, displays all volume names.

– Specifies that you want to read the names of one or more volumes to list from stdin (for example, – << /work/vol_list.txt).

Description

Prerequisite: You must be logged in to the engine hosting the master Metadata server to run this command.

When you specify multiple parameters in this command, these behaviors apply:

- Flags work as AND coordinators.
- Objects work as OR coordinators.
- When flags and objects are combined, the they are listed in this manner: WHERE flag_1 AND flag_2 AND ... AND flag_n AND (object_1 OR object_2 OR ... OR object_m).
- When flags or objects are not required and not specified, this command searches all possible values unless a default overrides this rule.

For example, if you specify –pool, –state, and two volume names, the search would be performed in the following manner: WHERE pool AND state AND (volume_name_1 OR volume_name_2).

If you do not specify a listing parameter, this command displays the following default information for each volume:

- The name of the volume.
- The state of the volume.
- The size of the volume, in MB.

  Note: When the Metadata server creates a volume, it uses some space for its internal bookkeeping (including the disk label). This command reports the size of the volume as the amount of space that can be used to store data, which does not include the space used for bookkeeping.
- The amount of space being used, in MB.
- The percent of space being used.

If you specify the **–l** listing parameter, this command displays the following information in addition to the default information for each volume:

- The operating-system device path.
- The SAN File System volume ID.
- The volume description.

Examples

Display a list of volumes The following example lists all the attributes of all the volumes:

```
tanktool> lsvol –l

Name   State           Pool        Size (MB)   Used (MB)   Used (%)
=====================================================================
Vol1   Activated       P_Accting         700         350         50
Vol2   Being Removed   P_Accting         700         175         25

Device Path     Vol ID   Description
====================================
/dev/hrvpatha   912478   Department 1
/dev/hrvpathb   807890   Department 2
```

Related topics

- “activatevol” on page 163
- “chvol” on page 178
- “mkvol” on page 240
- “reportvolfiles” on page 245
- “rmvol” on page 258
- “suspendvol” on page 294

mkcontainer

Creates a new fileset (container).

Syntax

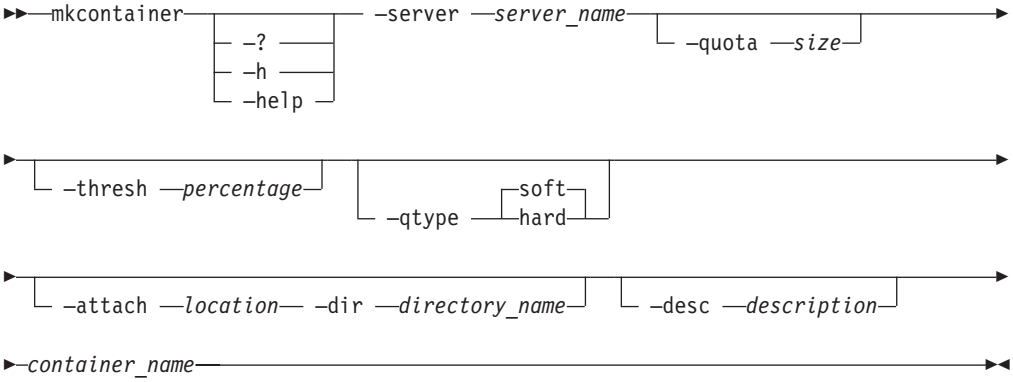

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–server *server_name*
Specifies the Metadata server to host the fileset. This name can be up to 32 characters in length.

Note: This Metadata server can be in any state as long as it is part of the cluster.

–quota *size*
Specifies the maximum size limit, in megabytes, for the specified fileset that, when exceeded, causes the Metadata server to generate an alert. You can specify a number from 0 to 1 073 741 824 MB (1 PB). The default size is 0.

If not specified or set to 0, there is no quota size limit for this fileset, and alerts are not sent.

–thresh *percentage*
Specifies the maximum percentage (alert threshold) of the specified quota size for this fileset that, when exceeded, causes the Metadata server to generate an alert. You can specify a value between 0 and 100. The default alert percentage is 80.

If the **–thresh** parameter is set to 0 or if the quota size is not specified, no alerts are generated.

–qtype hard | soft
Specifies the quota type for the fileset. You can specify one of the following values:

hard A hard quota produces a log message and potential alert when the quota is met, and denies requests for more space.

soft A soft quota produces a log message and potential alert when the quota size is exceeded, but grants requests for more space.

This is the default type.

–attach *location*
Specifies the existing directory path (relative to the root of the global namespace and without the / prefix) at which to attach the fileset. This directory ***must*** already exist.

The root of the global namespace must be included in the directory path. For example, to attach the fileset to the /homes location, specify –attach sanfs/homes. To attach the fileset directly to the root of the global namespace, use –attach sanfs.

Use only forward slashes (/) in the directory path for delimiters.

Note: If the **–attach** parameter is not specified, the fileset will not be attached to the global namespace.

–dir *directory_name*
Specifies the directory name for the root of the fileset directory tree. The directory name ***must not*** already exist. This directory can be up to 256 characters in length.

The directory must not contain backslash (\) or forward slash (/) characters as delimiters.

The attach point (*location*/*directory_name*) appears as a drive and directory on Windows clients or as a file system on UNIX-based clients. It is recommended that you map the name of the directory to its associated fileset by using the same name as the fileset plus an additional suffix to distinguish it as the directory and not the fileset. For example, if you want to attach the fileset named work to the root of the global namespace, you would specify "sanfs" for the *location* and specify "work_ap" for the *directory_name*. On a Windows client, the "work_ap" directory would appear under the "sanfs" drive.

–desc *description*
Specifies a description for the fileset. The description must be enclosed in matching single (') or double (") quotation marks if it contains any blank characters. This description can be up to 256 characters in length.

The default is an empty string.

container_name
Specifies the name to assign to the new fileset (or container). This name can be up to 256 characters in length.

Description

Prerequisites:

1. You must have Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

Note: You must attach the fileset to the global namespace using the **–attach** and **–dir** parameters in this command or the **attachcontainer** command before clients can access the fileset.

Newly created filesets have permissions set to the following:

- 555 permissions when viewed from UNIX-based clients.
- Read control (to Everyone) and Write DAC/Owner (for administrator or root clients) from Windows-based clients.

Examples

Create filesets The following example creates a fileset named cnt_A, assigns the fileset to the ST1 Metadata server, sets the quota to 1 000 MB, specifies to send alerts when the quota reaches 70%, and attaches the fileset to the sanfs/homes directory.

```
tanktool> mkcontainer –server ST1 –quota 1000 –thresh 70 –qtype soft
–attach sanfs/homes –dir cntA_ap cnt_A
Container "cnt_A" created.
```

Related topics

- “attachcontainer” on page 166
- “chcontainer” on page 174
- “detachcontainer” on page 183
- “lscontainer” on page 198
- “rmcontainer” on page 251

mkdrfile

Creates a new system-metadata disaster-recovery dump file.

Syntax

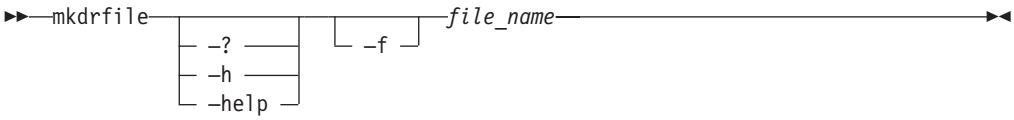

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–f Overwrites the system-metadata disaster-recovery dump file if a file with the specified name already exists.

file_name
Specifies the name of the new system-metadata disaster-recovery dump file. This name can be up to 250 characters in length.

Description

Prerequisites:

1. You must have Backup, Operator, or Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

This command is used in conjunction with the **builddrscript** command to create a set of recovery scripts that are used to recreate the system metadata in the event of a disaster. The **builddrscript** command is run against the system-metadata disaster-recovery file created by the **mkdrfile** command.

Examples

Create a dump file The following example creates a disaster-recovery dump file named “dr1.”

```
tanktool> mkdrfile dr1
Disaster recovery file "dr1" was created successfully.
```

Related topics

- “builddrscript” on page 168
- “lsdrfile” on page 202
- “rmdrfile” on page 252

mkimage

Creates a near-instantaneous FlashCopy image of the file layout and contents of the specified fileset and stores it in that fileset.

Syntax

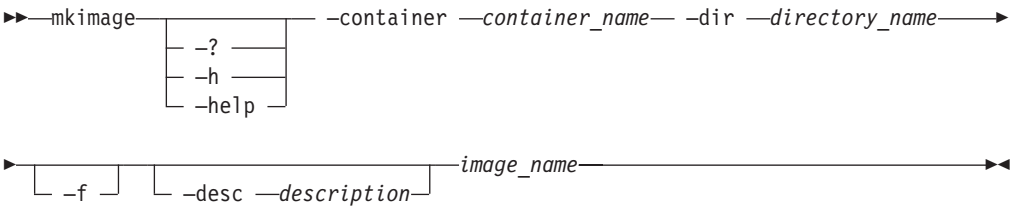

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–container *container_name*
Specifies the name of the fileset for which to create a FlashCopy image (or container).

–dir *directory_name*
Specifies the file-system directory under which you want to access the FlashCopy image. This directory can be up to 256 characters in length and must not contain backslash (\) or forward slash (/) characters as delimiters.

The directory name is created under the .flashcopy directory in the fileset. The full path of any file in the FlashCopy image is under the directory *attach_point*/.flashcopy/*directory*, where *attach_point* is the location to which the fileset attaches in the global namespace (for example, sanfs/cnt_A/.flashcopy/fc1dir).

Note: The .flashcopy directory is a hidden directory. Windows-based clients must use special operating-system commands to view this directory.

–f Forces the Metadata server to delete the oldest FlashCopy image if the fileset already contains the maximum number of FlashCopy images. The oldest image is deleted even if the image is currently in use.

Attention: Any open files associated with the oldest image will be lost.

–desc *description*
Specifies a description for the FlashCopy image. The description must be enclosed in matching single (') or double (") quotation marks if it contains any blank characters. This description can be up to 256 characters in length.

The default is an empty string.

image_name
Specifies the name of the FlashCopy image to create. This name must be unique within the fileset and can be up to 256 characters in length.

Note: This name must not contain colons (:).

Description

Prerequisites:

1. You must have Backup, Operator, or Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

This command does not create FlashCopy images of nested filesets. You must create a FlashCopy image for each fileset in the hierarchy.

After you create a FlashCopy image, you cannot modify the name, description, or directory.

Until the **mkimage** command completes, you can view files in the fileset but not modify them.

You cannot use the **mkimage** command while the **reverttoimage** command is in progress.

A fileset is limited to 32 FlashCopy images at one time.

Examples

Create a FlashCopy image The following example creates a FlashCopy image for the fileset cnt_A.

```
tanktool> mkimage –container cnt_A –dir website cntA_image
Image "cntA_image" successfully created.
```

Related topics

- “lsimage” on page 207
- “reverttoimage” on page 249
- “rmimage” on page 253

mkpolicy

Creates a policy.

Syntax

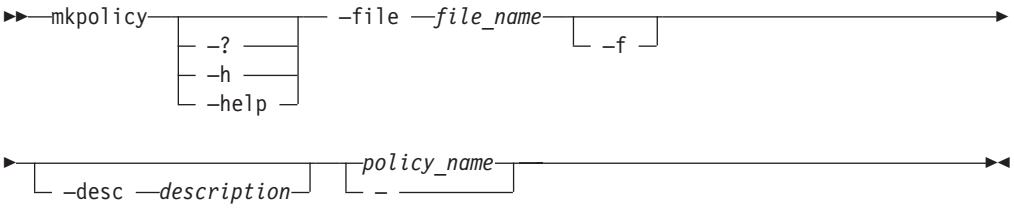

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–file *file_name*
Specifies the fully qualified path and name of the file that contains the set of rules to be included in the new policy. The maximum size of this file is 8 KB.

Use the following format to add rules to this file:

```
VERSION 1

rule 'stgRule1' set stgpool 'pool1' for fileset ('cnt_A')
rule 'stgRule2' set stgpool 'pool2' where NAME like '%.doc'
rule 'stgRule3' set stgpool 'pool3' where DAYOFWEEK(CREATION_DATE) == 1
rule 'stgRule4' set stgpool 'pool4' where USER_ID <= 100
```

–f Forces the Metadata server to overwrite an existing policy with the same name as the specified policy.

–desc *description*
Specifies a description for the policy. The description must be enclosed in matching single (') or double (") quotation marks if it contains any blank characters. This description can be up to 256 characters in length.

The default is an empty string.

policy_name
Specifies the names of the policy to create. This name can be up to 256 characters in length.

– Specifies that you want to read the names of the policy to create from stdin (for example, – << /work/policies_list.txt).

Description

Prerequisites:

1. You must have Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

Each policy can be up to 32 KB.

To update the policy rules or attributes through the Administrative command-line interface, you must create a new policy.

Examples

Create a policy The following example creates a policy (*test_policy*) using the rules in the */tmp/my_rules.txt* file:

```
tanktool> mkpolicy –file /tmp/my_rules.txt –desc "Test Policy" test_policy
Policy test_policy created.
```

Related topics

- "catpolicy" on page 171
- "lspolicy" on page 213
- "rmpolicy" on page 255
- "usepolicy" on page 296

mkpool

Defines one or more new user storage pools.

Syntax

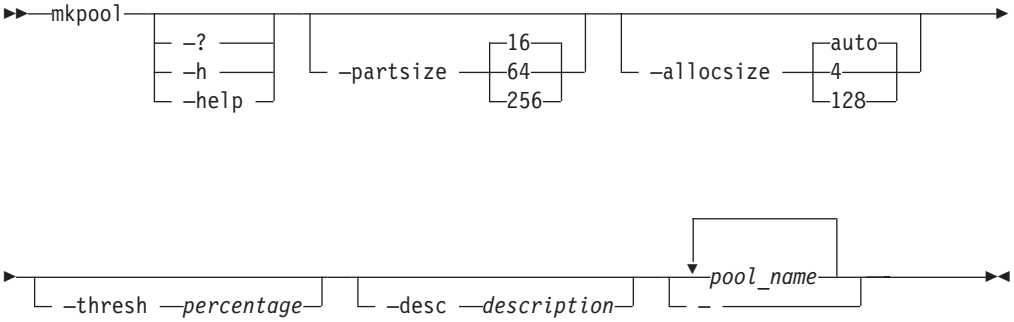

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–partsize *size*
Specifies the partition size, in megabytes, to use when a fileset allocates space. The partition size can be set to 16, 64, or 256 MB. The default size is 16 MB.

Attention: You cannot change the partition size after it is set.

–allocsize *size*
Specifies the allocation size, in kilobytes, by which a file that resides on a storage pool is to be extended. The allocation size can be set to auto, 4, or 128. If you specify auto, the system sets the size automatically. The default value is auto.

Attention: You cannot change the block size after it is set.

–thresh *percentage*
Specifies a percentage of the storage pool's estimated capacity that, when reached or exceeded, causes the Metadata server to generate an alert. If the capacity drops and then reaches the percentage again, additional alerts are sent.

You can specify a value between 0 and 100. The default alert threshold percentage is 80. If set to 0, no alerts are generated.

–desc *description*
Specifies a description for the storage pool. The description must be enclosed in matching single (') or double (") quotation marks if it contains any blank characters. This description can be up to 256 characters in length.

The default is an empty string.

pool_name
Specifies the names of one or more storage pools to create. Each name can be up to 256 characters in length.

– Specifies that you want to read the names of one or more storage pools to create from stdin (for example, – << /work/stgpools_list.txt).

Description

Prerequisites:

1. You must have Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

You can create multiple storage pools with the same values for the partition size, allocation size, and alert threshold in a single command; however, you cannot set the name or description for multiple storage pools.

Examples

Create two storage pools The following example creates two storage pools (stgpool1 and stgpool2), sets the partition size on both to 256 MB, and specifies to send alerts when the allocated capacity reaches 70%:

```
tanktool> mkpool –partsize 256 –thresh 70 stgpool1 stgpool2
Storage pool stgpool1 created.
Storage pool stgpool2 created.
```

Related topics

- “chpool” on page 176
- “lspool” on page 216
- “rmpool” on page 256
- “setdefaultpool” on page 262

mkvol

Adds one or more volumes to a storage pool.

Syntax

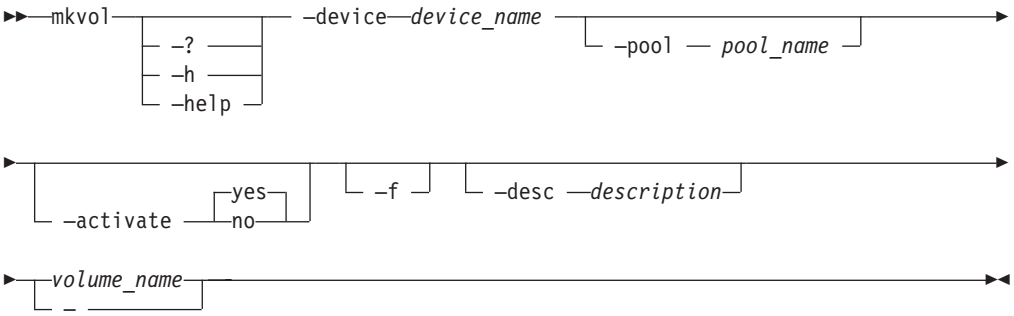

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–device *device_name*
Specifies the name of the local device to add to the specified storage pool (for example, */dev/rsdc* on Linux).

–pool *pool_name*
Specifies the name of the storage pool to which to add the new volumes. The storage pool is either a user-defined, default, or system storage pool. If not specified, this command adds the volume to the default storage pool.

–activate
Specifies whether the volume is activated. The Metadata server can allocate new data only on activated volumes. Possible values are:

yes Specifies that the Metadata server can allocate new data on the volume. This is the default value.

no Specifies that the Metadata server cannot allocate new data on the volume.

–f Forces the Metadata server to add the volume and write a new label to the volume if the volume already has a valid SAN File System label.

Note: You can use **–f** only if the volume is not assigned to another storage pool in the same cluster.

–desc *description*
Specifies the description for the volume. The description must be enclosed in matching single (') or double (") quotation marks if it contains any blank characters. This description can be up to 256 characters in length.

The default is an empty string. Specifying this parameter with no description results in a syntax error.

volume_name
Specifies the names of one or more volumes to add. This name must be unique within the fileset and can be up to 256 characters in length.

– Specifies that you want to read the names of one or more volumes to add from stdin (for example, – << /work/vol_list.txt).

Description

Prerequisites:

1. You must have Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.
3. You must format the applicable LUNs before invoking this command.

Examples

Create a volume The following example creates a volume and adds it to the storage pool named fast:

```
tanktool> mkvol –device /dev/rvpatha –pool fast –desc "fast vol 1" first_vol
Volume first_vol created and added to pool fast.
```

Related topics

- "activatevol" on page 163
- "chvol" on page 178
- "lsvol" on page 227
- "reportvolfiles" on page 245
- "rmvol" on page 258
- "suspendvol" on page 294

quiescecluster

Changes the state of all Metadata servers in the cluster to one of three quiescent states.

Syntax

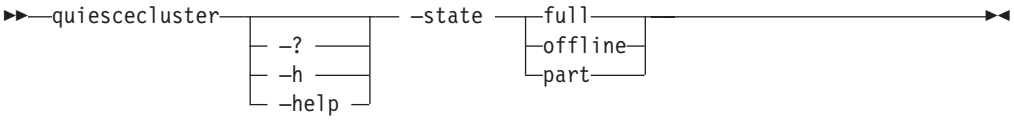

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–state
Specifies the state to which you want to change the cluster. You can specify one of these states:

full Changes the state to fully quiescent. This parameter flushes all client caches and long-running processes, resulting in no activity on the Metadata servers.

You would put the cluster in this state to produce self-consistent LUN-based backups for a known FlashCopy image.

offline
Changes the state to offline. This parameter terminates all current client sessions and prevents new client sessions from being started, but allows all Metadata server I/O to continue. The offline state is intended to be used as an administrative mode to, for example, reassign a fileset to another Metadata server.

You would put the cluster in this state to restrict clients from accessing the cluster.

part Changes the state to partly quiescent. This parameter prevents new client sessions from starting. Open files are still accessible, but not new files can be opened or created while in this state.

You would put the cluster in this state to produce *dirty* LUN-based backups without requiring clients and applications to stop, and to allow client applications that use direct I/O (such as database servers) to back up their data.

Description

Prerequisites:

1. You must have Operator or Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

The quiescent states restrict activity on all Metadata servers in the cluster. You would put the Metadata servers in this state in preparation for backup tasks that require the client sessions, file I/O, client metadata, or system metadata to be inactive.

Examples

Change the cluster state The following example changes the cluster state to fully quiescent.

```
tanktool> quiescecluster–state full
Cluster successfully in fully quiescent state.
```

Related topics

- "resumecluster" on page 248
- "startcluster" on page 272
- "statcluster" on page 278
- "stopcluster" on page 290

quit

Ends the tanktool session.

Syntax

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

Examples

Ends a tanktool session The following example ends the tanktool session:

```
tanktool> quit
shell>
```

Related topics

- "exit" on page 185

reportvolfiles

Displays a list of files (and their attributes) in the specified volume.

Syntax

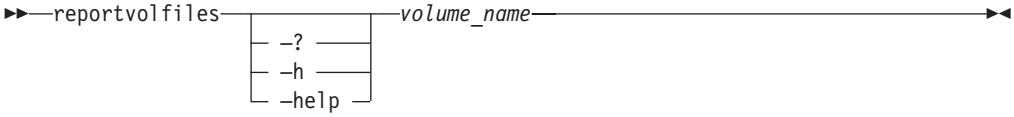

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

volume_name
Specifies the name of the volume to report.

Description

Prerequisites:

1. You must have Backup, Operator, or Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

Description

This command displays the following information for each file in the specified volume:

- Fileset name
- Fileset directory
- Relative file path
- File name

This information is displayed in the following format:

fileset:*fileset_directory*/*relative_file_path*/*file_name*

This command is useful for performing a file-based restore operation after a user volume failure because it can list the files in a failed user volume; however, it cannot list files on a failed system volume.

Examples

List files in a volume The following example displays a list of files in volume vol2.

```
tanktool> reportvolfiles vol2
Root:Work/doc/language.doc
NotesR5D:InfoDev/IDPlan.doc
Resources:Personnel/dept/manager/people.txt
```

Related topics

- "lsvol" on page 227
- "lscontainer" on page 198

resetadmuser

Forces all administrative users to log in again.

Syntax

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

Description

Prerequisite: You must have Administrator privileges to use the command.

Note: This command operates only against the local engine.

To increase Common Information Model (CIM) performance, the Administrative agent retains a copy of the administrative user's password after the LDAP server validates it. This copy is used for a certain period of time to avoid the overhead of going to the LDAP server to authenticate the user for every CIM request. If a user's password has been changed in the LDAP server, use this command to ensure that the user must immediately use the new password rather than continuing to use the Administrative agent's copy.

Examples

Reset the Administrative Users' Passwords The following example resets all administrative user's passwords in the CIM cache:

```
tanktool> resetadmuser
Users reset.
```

Related topics

- "lsadmuser" on page 187

restartengine

Reboots one or more storage engines.

Syntax

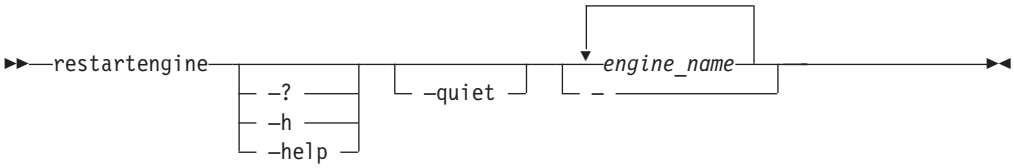

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–quiet
Turns off confirmation prompts for this command.

engine_name
Specifies the names of one or more storage engines to restart.

– Specifies that you want to read the names of one or more storage engines to restart from stdin (for example, – << /work/engine_list.txt).

Description

Prerequisite: You must have Operator or Administrator privileges to use the command.

You can stop and restart the local engine, but you cannot start the local engine.

Examples

Restart engines The following example restarts engine ST1.

```
tanktool> restartengine ST1
Are you sure you want to restart engine ST1? [y/n] y
Engine ST1 restarted successfully.
```

Related topics

- “lsengine” on page 204
- “startengine” on page 273
- “statengine” on page 282
- “stopengine” on page 291

resumecluster

Brings all Metadata servers in the cluster to the online state.

Syntax

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

Description

Prerequisites:

1. You must have Operator or Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

Examples

Resume the cluster The following example resumes normal activity on the cluster.

```
tanktool> resumecluster
Cluster successfully returned to online state.
```

Related topics

- “quiescecluster” on page 242
- “startcluster” on page 272
- “statcluster” on page 278
- “stopcluster” on page 290

reverttoimage

Reverts the current fileset to a specified FlashCopy® image of the file layout and contents.

Syntax

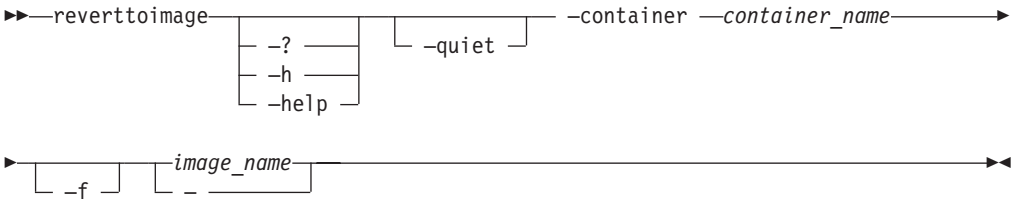

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–quiet
Turns off the prompt to confirm the revert operation.

–container *container_name*
Specifies the name of the fileset (or container) that you want to revert to a previous FlashCopy image.

–f Forces the Metadata server to revert to the specified FlashCopy image if a client is reading files in the current image. Changes made to files in the fileset since the specified FlashCopy image was created will be lost.

image_name
Specifies the name of the FlashCopy images to which you want to revert.

– Specifies that you want to read the name of the FlashCopy image to which you want to revert from stdin (for example, – << /work/image_list.txt).

Description

Attention: When you revert a fileset to a specified target FlashCopy image, the target FlashCopy image and all FlashCopy images taken between the current fileset and target FlashCopy image are deleted. The target FlashCopy image becomes the primary image for the fileset and no longer appears as an image listed in the .flashcopy directory.

Prerequisites:

1. You must have Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

Unless you specify the **–f** parameter, there ***must not*** be any client activity in the FlashCopy images being deleted or in the current fileset image.

You cannot revert to a FlashCopy image when nested filesets exist within the fileset. You must manually detach the nested filesets before running the **reverttoimage** command.

The .flashcopy directory is a hidden directory. Windows-based clients must use special operating-system commands to view this directory.

Depending on the age of the specified FlashCopy image and the amount of unique file data in the image tree, the revert operation could result in significant background activity to clean up the file-system objects that are no longer referenced.

Until the **reverttoimage** command completes, you cannot perform any activity on the fileset except the **rmimage** and **reverttoimage** commands. If you issue a second **reverttoimage** command before the first one completes, you can only revert to a FlashCopy image that existed before the FlashCopy image being reverted by the first command.

Because the specified FlashCopy image is deleted after you issue the **reverttoimage** command, it is recommended that you preserve a secondary backup of the fileset before using the command for future use or disaster recovery.

To preserve the fileset hierarchy, you must reattach any nested filesets that you detached before issuing the command when this command completes. It is recommended that you note the hierarchy using the **lscontainer** command before reverting to a FlashCopy image.

Examples

Revert to a previous FlashCopy image The following example reverts the fileset *cnt_A* to the FlashCopy image *cntA_image*.

```
tanktool> reverttoimage –container cnt_A cntA_image
Are you sure you want to revert to FlashCopy image cntA_image? y/n y
Image "cntA_image" successfully restored.
```

Related topics

- “detachcontainer” on page 183
- “lscontainer” on page 198
- “lsimage” on page 207
- “mkimage” on page 234
- “rmimage” on page 253

rmcontainer

Removes one or more empty, detached filesets (containers) and optionally the files in the filesets, including any FlashCopy® images.

Syntax

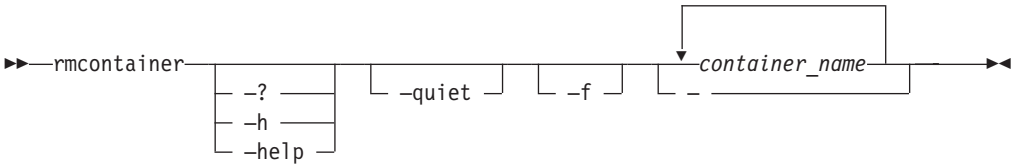

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–quiet
Turns off the prompt to confirm the remove operation.

–f Forces the Metadata server to delete any files and directories before removing the filesets.

container_name
Specifies the names of one or more filesets (or containers) to remove.

– Specifies that you want to read the names of one or more filesets to remove from stdin (for example, – << /work/cnt_list.txt).

Description

Prerequisites:

1. You must have Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

Unless you specify the **–f** parameter, the fileset ***must*** be detached and empty before it can be deleted.

You cannot remove the ROOT fileset or a fileset that is reference by the active policy.

Examples

Remove a fileset The following example removes the fileset named cnt_A.

```
tanktool> rmcontainer cnt_A
Are you sure you want to delete container cnt_A? [y/n] y
Container cnt_a removed.
```

Related topics

- “attachcontainer” on page 166
- “chcontainer” on page 174
- “detachcontainer” on page 183
- “lscontainer” on page 198
- “mkcontainer” on page 230

rmdrfile

Deletes an existing system-metadata disaster-recovery dump file.

Syntax

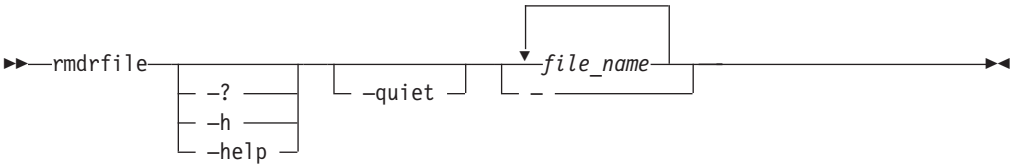

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–quiet
Turns off the prompt to confirm the change operation

file_name
Specifies the names of one or more system-metadata disaster-recovery dump files to remove.

– Specifies that you want to read the names of one or more system-metadata disaster-recovery dump files to remove from stdin.

Description

Prerequisite: You must have Backup, Operator, or Administrator privileges to use the command.

Note: This command operates only against the local engine.

Examples

Remove a dump file The following example removes the dump file named dr1.

```
tanktool> rmdrfile dr1
Disaster recovery file "dr1" removed.
```

Related topics

- "builddrscript" on page 168
- "lsdrfile" on page 202
- "mkdrfile" on page 233

rmimage

Deletes one or more FlashCopy® images for a specific fileset.

Syntax

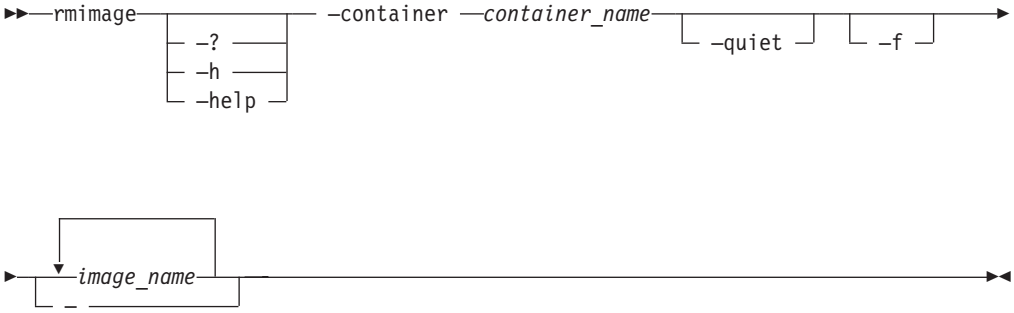

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–container *container_name*
Specifies the name of the fileset (or container from which to delete one or more FlashCopy images.

–quiet
Turns off the prompt to confirm the delete operation.

–f Forces the Metadata server to delete a FlashCopy image if one or more files are open in the image tree. Changes that were made to the open files but not saved will be lost.

image_name
Specifies the names of one or more FlashCopy images to delete.

– Specifies that you want to read the names of one or more FlashCopy images to delete from stdin (for example, – << /work/image_list.txt).

Description

Prerequisites:

1. You must have Backup, Operator, or Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

Unless you specify the **–f** parameter, you cannot delete a FlashCopy image if there is client activity on that image.

Depending on the age of the FlashCopy image and the amount of unique file data in the image tree, the delete operation might result in significant background activity to clean up the file-system objects that are no longer referenced.

Examples

Delete a FlashCopy image The following example deletes the FlashCopy image (*CA_image*) from the fileset *cnt_A*.

```
tanktool> rmimage –container cnt_A CA_image
Are you sure you want to delete FlashCopy image CA_image? [y/n] y
Image "CA_image" successfully deleted.
```

Related topics

- “lsimage” on page 207
- “mkimage” on page 234
- “reverttoimage” on page 249

rmpolicy

Deletes one or more inactive policies.

Syntax

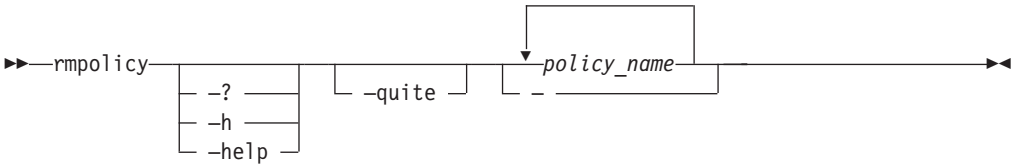

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–quiet
Turns off the prompt to confirm the remove operation.

policy_name
Specifies the names of one or more inactive policies to delete.

– Specifies that you want to read the names of one or more inactive policies to delete from stdin (for example, – << /work/policies_list.txt).

Description

Prerequisites:

1. You must have Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

Examples

Remove a policy The following example removes the policy named *test_policy*.

```
tanktool> rmpolicy test_policy
Are you sure you want to delete policy test_policy? [y/n] y
Policy test_policy removed.
```

Related topics

- “catpolicy” on page 171
- “lspolicy” on page 213
- “mkpolicy” on page 236
- “usepolicy” on page 296

rmpool

Deletes one or more empty, unreferenced storage pools.

Syntax

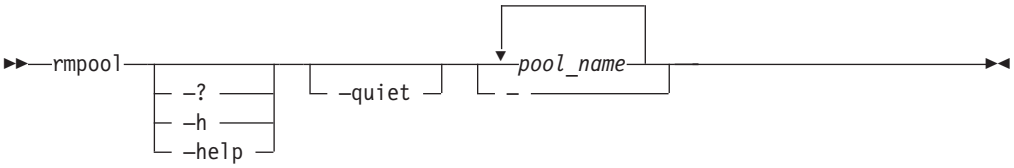

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–quiet
Turns off the prompt to confirm the delete operation.

pool_name
Specifies the names of one or more storage pools to delete.

– Specifies that you want to read the names of one or more storage pools to delete from stdin (for example, – << /work/stgpools_list.txt).

Description

Prerequisites:

1. You must have Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

You cannot delete a storage pool that is not empty. You must remove all volumes from the storage pool before you can delete it.

You cannot delete a storage pool that is referenced by the active policy.

Examples

Remove storage pools The following example removes storage pool stgpool1.

```
tanktool> rmpool stgpool1
Are you sure you want to delete storage pool stgpool1?
Storage pool stgpool1 removed.
```

Related topics

- “chpool” on page 176
- “lspool” on page 216
- “mkpool” on page 238
- “setdefaultpool” on page 262

rmsnmpmgr

Removes an SNMP manager (recipient).

Syntax

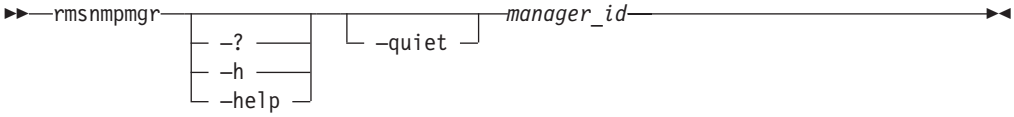

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–quiet
Turns off the prompt to confirm the delete operation.

manager_id
Specifies the number that identifies the SNMP manager (recipient) to delete. Because this number is not static, you must acquire the new number each time a change is made to the list of managers.

Description

Prerequisites:

1. You must have Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

If all SNMP managers are removed, no SNMP traps are generated.

When you change a disruptive cluster setting, dynamic cluster settings (such as SNMP settings) cannot be modified until you reboot the cluster.

Examples

Removes an SNMP manager The following example removes an SNMP manager.

```
tanktool> rmsnmpmgr 1
Are you sure you want to remove SNMP manager 1? [y/n] y
SNMP manager 1 successfully removed.
```

Related topics

- “addsnmpmgr” on page 164
- “catlog” on page 169
- “clearlog” on page 180
- “lssnmpmgr” on page 225
- “lstrapsetting” on page 226
- “settrap” on page 268
- “SNMP” on page 18

rmvol

Removes one or more volumes from a storage pool and redistributes the contents to other volumes in the same storage pool.

Syntax

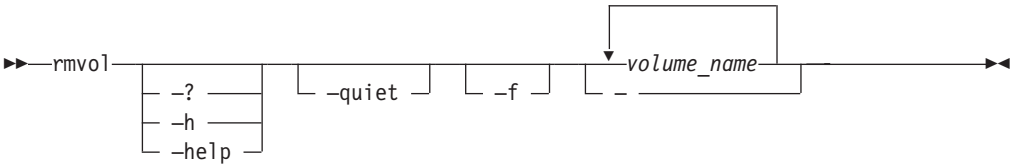

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–quiet
Turns off the prompt to confirm the delete operation.

–f Removes all files in the specified volumes and disables draining the files to another volume in the storage pool. If there is a file in the volume, specifying this parameter causes the entire file to be deleted, even if parts of the file reside on multiple volumes.

The **–f** parameter is not allowed when you remove volumes from the system storage pool.

Note: There is no automatic-recovery process when you specify this parameter. You can look at the activity log to get a list of the files for which failures occurred, and then restore those files manually. You can also use the **reportvolfiles** command before invoking the **rmvol** command and review the file list to make sure that the files are not needed or are archived.

volume_name
Specifies the names of one or more volumes to delete.

– Specifies that you want to read the names of one or more volumes to delete from stdin (for example, – << /work/vol_list.txt).

Description

Prerequisites:

1. You must have Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

You do not need to deactivate a volume before you remove it.

Before removing a volume, SAN File System moves (drains) the contents of the volume across other available volumes in the same storage pool. If the storage pool does not have sufficient space available in other volumes to move all of the data contained in the specified volume, the removal fails and the Metadata server suspends the volume (the Metadata server cannot allocate new data on that volume).

A volume must be empty to be removed from a storage pool. If one or more files cannot be moved to another volume in the same storage pool, you can specify the **–f** parameter to discard those files. If you do not specify the **–f** parameter and the files cannot be removed, the removal fails.

If an error occurs during volume removal, the volume remains in a suspended state so new allocations will not be satisfied from the volume being removed. To activate the volume, use the **activatevol** command.

You must explicitly remove a volume if you want to reuse it. Removing a volume removes the label and frees the device for reuse. If you want to add a device that is in the in-use state, you must remove it using the **–f** option before you can add it as a volume.

Examples

Remove two volumes The following example removes volumes vol1 and vol2.

```
tanktool> rmvol vol1 vol2
Are you sure you want to delete volume vol1? [y/n] y
Volume vol1 removed.
Are you sure you want to delete volume vol21? [y/n] y
Volume vol2 removed.
```

Related topics

- “activatevol” on page 163
- “chvol” on page 178
- “lsvol” on page 227
- “mkvol” on page 240
- “reportvolfiles” on page 245
- “suspendvol” on page 294

setcontainerserver

Reassigns an existing fileset (container) to be hosted by a different Metadata server.

Syntax

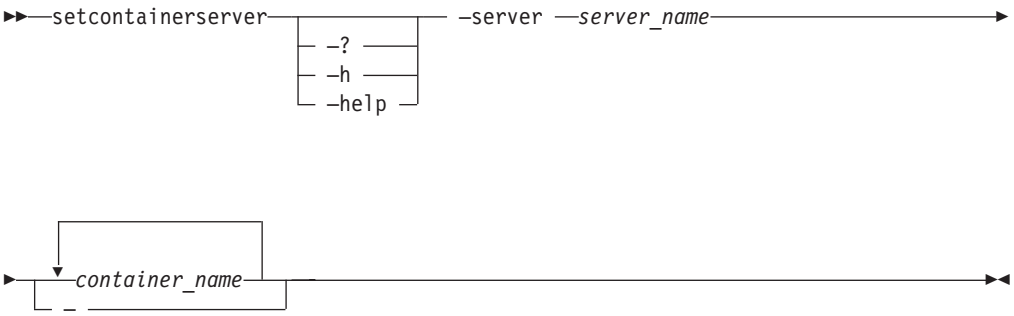

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–server *server_name*
Specifies the new Metadata server to host the fileset.

Note: If you specify this parameter, the Metadata server currently hosting the fileset must be either offline or not running, and the new host Metadata server must be part of the cluster. The new host Metadata server must be in the online, offline, or not running state.

container_name
Specifies the names of one or more filesets (or containers) to reassign to a new Metadata server.

–
Specifies that you want to read the names of one or more filesets to reassign from stdin (for example, – << /work/cnt_list.txt).

Description

Prerequisites:

1. You must have Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.
3. The current Metadata server must be in one of the following states:
   - Not running
   - Offline
4. The target Metadata server must be part of the cluster and must be in one of the following states:
   - Not running
   - Offline
   - Online

If current Metadata server is online and the target Metadata server is either online or offline, then the cluster is set to the offline state (administrative mode). If

current Metadata server is offline, and the target Metadata server is either online or offline, then the state of the cluster does not change.

Examples

Reassign a fileset The following example reassigns two filesets (*cnt_A* and *cnt_B*) to Metadata server ST1.

```
tanktool> setcontainerserver –server ST1 cnt_A cnt_B
Container cnt_A assigned to server ST1.
Container cnt_B assigned to server ST1.
```

Related topics

- "attachcontainer" on page 166
- "chcontainer" on page 174
- "detachcontainer" on page 183
- "lscontainer" on page 198
- "mkcontainer" on page 230
- "rmcontainer" on page 251

setdefaultpool

Changes a user storage pool to the default storage pool, and changes the previous default storage pool to a regular, nondefault user storage pool.

Syntax

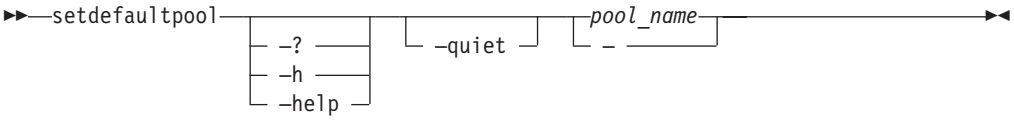

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–quiet
Turns off the prompt to confirm the change operation.

pool_name
Specifies the name of the storage pool to set as the default.

– Specifies that you want to read the name of the storage pool to set as the default from stdin (for example, – << /work/stgpool_list.txt).

Description

Prerequisites:

1. You must have Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

You cannot change the SYSTEM storage pool using this command.

Examples

Set the DEFAULT storage pool The following example changes the storage pool named default_stgpool to be the new DEFAULT storage pool:

```
tanktool> setdefaultpool default_stgpool
Are you sure you want to set default_stgpool as the default pool? [y/n] Y
Pool default_stgpool is now the DEFAULT pool.
```

Related topics

- "chvol" on page 178
- "lsvol" on page 227
- "mkpool" on page 238
- "rmvol" on page 258

setmaster

Sets a subordinate Metadata server as the new master Metadata server in the event of an irrecoverable loss of the current master Metadata server.

Syntax

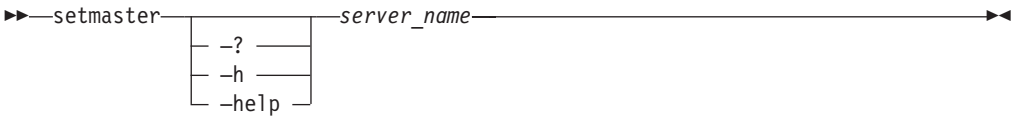

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

server_name
Specifies the name of the subordinate Metadata server to set as the master Metadata server.

Description

Prerequisites:

1. You must have Administrator privileges to use the command.
2. You must be logged in to the engine hosting the *subordinate* Metadata server to run this command.

Attention: When you invoke the **setmaster** command, applications will notice a pause in service, which could adversely affect application that are sensitive to timeouts.

Attention: If the master Metadata server is not shut down, you could destroy all of your data. Before using the setmaster command, perform the following steps to verify that the master Metadata server is not shut down:

1. From the engine hosting the master Metadata server, perform the following steps:
   a. Verify that the master Metadata server is *offline* by invoking the **lsserver** command.
   b. If the master Metadata server is *online*, invoke the **stopserver** command to stop the master Metadata server.

      Do not stop the subordinate Metadata servers.
   c. Verify again that the master Metadata server is *offline* by invoking the **lsserver** and **statcluster –netconfig** commands.
2. Verify that all other Metadata servers are subordinates by invoking the **lsserver** command from each engine hosting the remaining Metadata servers.
3. From the engine hosting the subordinate Metadata server that you want to become the master, perform the following steps:
   a. Verify that the engine hosting the master Metadata server is shut down by invoking the **lsengine** command.

      If the engine is running, use the **stopengine** or **stopengine –f** command to shut it down.

      Note: If you cannot shut down the engine remotely or through the Administrative CLI, manually power off the engine.
   b. Verify again that the engine hosting the master Metadata server is shut down by invoking the **lsengine** command.

Before stopping the master Metadata server, you should move the filesets that are being served by the master Metadata server to another Metadata server.

Setting a new master Metadata server may affects the clients. If the clients mount to the global namespace using the IP address of the master Metadata server, this address will become unavailable if the client reboots. Therefore, you must reset the IP address to which each client mounts. For AIX client, edit the /usr/tank/client/config/stclient.conf file, and set the server_hot_name variable.

The **setmaster** command demotes the master Metadata server to a subordinate. It does not drop the Metadata server from the cluster. When you restart the engine, the Metadata server joins the cluster as subordinate.

This command is useful when you want to replace hardware components of or upgrade the engine hosting the master Metadata server.

Examples

Set the master Metadata server The following example verifies that the master Metadata server ST0 is shut down and then sets the subordinate Metadata server ST1 as the master Metadata server.

From Metadata server ST0:

```
tanktool> lsserver
Name  State   Server Role  Containers  Last Boot
============================================================
ST0   Online  Master                0  Jul 10, 2003 7:00:17 AM
ST1   Online  Subordinate           2  Jul 10, 2003 7:00:24 AM

tanktool> stopservre ST0
```

```
Server "ST0" stopped gracefully.

tanktool> lsserver
Name  State        Server Role  Containers  Last Boot
=================================================================
ST0   Not Running  Master                -  -
ST1   Joining      Subordinate           2  Jul 10, 2003 7:00:24 AM

tanktool> statcluster –netconfig
Could not connect to the server. Please start the server.
```

From Metadata server ST1:

```
tanktool> lsserver
Name  State    Server Role  Containers  Last Boot
=============================================================
ST1   Joining  Subordinate           2  Jul 10, 2003 7:00:24 AM

tanktool> statcluster –netconfig
statcluster must be issued from the admin master server.

tanktool> stopengine ST0
Are you sure you want to stop engine "ST0"? [y/n]  y
Engine "ST0" was shutdown successfully.

tanktool> lsengine
Name  Engine IP      Boot State         Temp    Fans    Voltage
=============================================================
ST1   555.168.10.25  In OS              Normal  Normal  Normal
ST0   555.168.10.24  Unknown/Power Off  -       -       -

tanktool> setmaster ST1
Server "ST1" is the new master server.

tanktool> lsserver
Name  State    Server Role  Containers  Last Boot
=============================================================
ST1   Online   Master               2  Jul 10, 2003 7:05:12 AM
ST0   Unknown  Subordinate           2  Jul 10, 2003 7:05:12 AM
```

Related topics

- “lsengine” on page 204
- “lsserver” on page 221
- “Metadata server” on page 8
- “startserver” on page 276
- “stopcluster” on page 290
- “stopengine” on page 291

setoutput

Sets the output format for the Administrative command line interface.

Syntax

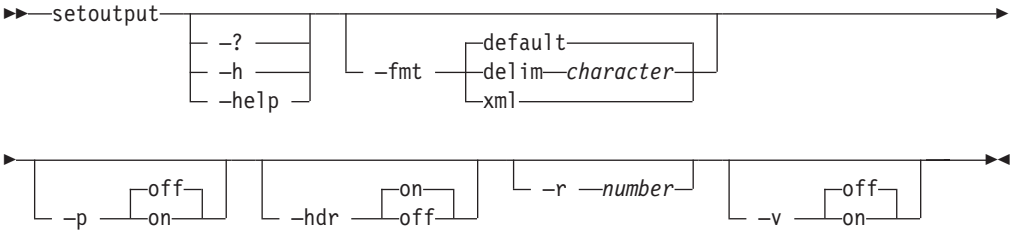

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–fmt
Specifies the format of the output. You can specify one of the following values:

default
Specifies to display output in a tabular format using spaces as the delimiter between the columns. This is the default value. For example:

```
Name      Type     Size (GB)   Used (GB)   Used (%)   Alert (%)
============================================================
DEFAULT   Default      10000        2500         25          80

Volumes   Partition Size (MB)   Description
=================================================
     10                    64   Default Storage Pool
```

delim *character*
Specifies to display output in a tabular format using the specified character to separate the columns. If you use a shell metacharacter (for example, * or \t) as the delimiting character, enclose the character in single quotation mark (′) or double quotation mark (″). A blank space is not a valid character. For example:

```
DEFAULT,Default,10000,2500,25,80,10,64,Default Storage Pool
```

xml Specifies to display output using XML format, for example:

```
<IRETURNVALUE>
<INSTANCE CLASSNAME="STC_StoragePool">
<PROPERTY NAME="Name" TYPE="string"><VALUE>DEFAULT_POOL</VALUE>
</PROPERTY>
<PROPERTY NAME="PoolType" TYPE="uint32"><VALUE>1</VALUE>
</PROPERTY>
<PROPERTY NAME="PartitionSize" TYPE="uint64"><VALUE>16</VALUE>
</PROPERTY>
<PROPERTY NAME="AlertPercentage" TYPE="uint16"><VALUE>80</VALUE>
</PROPERTY>
<PROPERTY NAME="Size" TYPE="uint64"><VALUE>0</VALUE></PROPERTY>
<PROPERTY NAME="SizeAllocated" TYPE="uint64"><VALUE>0</VALUE>
</PROPERTY>
<PROPERTY NAME="SizeAllocatedPercentage" TYPE="uint16"><VALUE>0
</VALUE></PROPERTY>
<PROPERTY NAME="NumberOfVolumes" TYPE="uint32"><VALUE>0</VALUE>
</PROPERTY>
```

```
<PROPERTY NAME="Description" TYPE="string"><VALUE>Default storage pool
</VALUE></PROPERTY>
</INSTANCE>
</IRETURNVALUE>
```

–p Specifies whether to display one page of text at a time or all text at once.

off Displays all text at one time. This is the default value when the **tanktool** command is run in single-shot mode.

on Displays one page of text at time. Pressing any key displays the next page. This is the default value when the **tanktool** command is run in interactive mode.

–hdr
Specifies whether to display the table header.

on Displays the table header. This is the default value.

off Does not display the table header.

–r *number*
Specifies the number of rows per page to display when the **–p** parameter is on. The default is 24 rows. You can specify a value from 1 to 100.

–v Specifies whether to enable verbose mode.

off Disables verbose mode. This is the default value.

on Enables verbose mode.

Description

The output format set by this command remains in effect for the duration of the tanktool session or until the options are reset either by using this command or by specifying a output-format parameter as part of a command.

Running this command with no parameters displays the current output settings in the default output format, for example:

```
Paging   Rows   Format    Header   Verbose
=========================================
off      -      default   on       off
```

Note: The output formats do not apply to help pages.

Examples

Set the output format The following example sets the output format to display in tabular form using a comma as the delimiter without header information

```
tanktool>setoutput -fmt delim , -hdr off
tanktool>lspool -l -type default
DEFAULT,Default,10000,2500,25,80,10,64,Default Storage Pool
```

settrap

Specifies whether an SNMP trap is generated and sent to all SNMP managers when a specific type of event occurs on the Metadata server.

Syntax

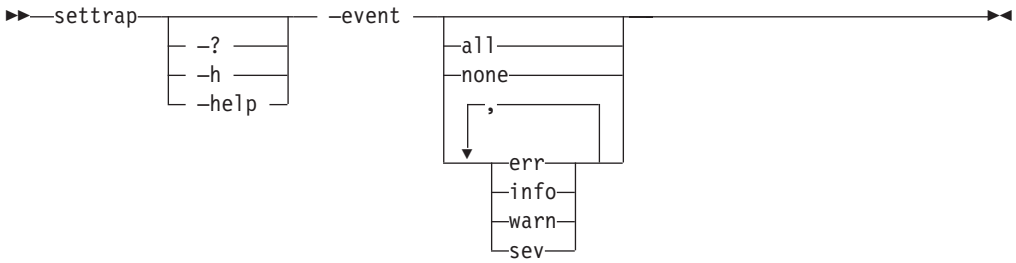

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–event
Specifies the events for which traps are to be sent. You can specify any of the following values:

all Generates an SNMP trap when an information, error, warning, or severe event occurs. This parameter cannot be combined with other values.

none Turns off SNMP traps for all events. This value cannot be combined with other values.

info Generates an SNMP trap when an information event occurs. This value can be combined with any value except all or none. Multiple values must be separated by a comma but no space.

err Generates an SNMP trap when an error event occurs. This value can be combined with any value except all or none. Multiple values must be separated by a comma but no space.

warn Generates an SNMP trap when a warning event occurs. This value can be combined with any value except all or none. Multiple values must be separated by a comma but no space.

sev Generates an SNMP trap when a severe event occurs. This value can be combined with any value except all or none. Multiple values must be separated by a comma but no space.

Description

Prerequisites:

1. You must have Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

Note: To generate traps, you must set the SNMP event types that you want to generate traps and you must add an SNMP manager (using the **addsnmpmgr** command).

SNMP traps of the specified type are generated for all SNMP managers.

When you change a disruptive cluster setting, dynamic cluster settings (such as SNMP settings) cannot be modified until you reboot the cluster.

Examples

Sets the SNMP traps The following example sets SNMP traps for severe and warning events.

```
tanktool> settrap –event sev,warn
SNMP trap event level was successfully set.
```

Related topics

- "addsnmpmgr" on page 164
- "catlog" on page 169
- "clearlog" on page 180
- "lssnmpmgr" on page 225
- "lstrapsetting" on page 226
- "rmsnmpmgr" on page 257

startautorestart

Enables the Metadata server to restart automatically if it is down.

Syntax

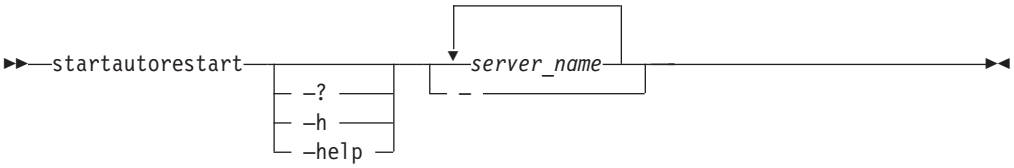

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

server_name
Specifies the names of one or more Metadata servers to enable them to restart automatically.

– Specifies that you want to read the names of one or more Metadata servers to enable to restart automatically from stdin (for example, – << /work/server_list.txt).

Description

Prerequisite: You must have Administrator privileges to use the command.

Note: If you run this command from an engine hosting a subordinate Metadata server, you can enable the Metadata server restart service on only the local Metadata server. If you run this command from the engine hosting the master Metadata server, you can enable the Metadata server restart service on any Metadata server.

The Metadata server restart service is enabled by default.

When a Metadata server is enabled to restart automatically, an SNMP trap is not sent when the Metadata server is restarted.

Manually stopping a Metadata server or cluster disables the Metadata server restart service for that Metadata server or cluster. Manually starting the Metadata server or cluster reenables the Metadata server restart service for that Metadata server or cluster.

If the Metadata server restart service attempts to restart a Metadata server and fails, the Metadata server restart service tries again to restart the Metadata server, up to the retry limit. After the retry limit, the Metadata server restart service is disabled.

Examples

Enable the Metadata server restart service The following example enables the Metadata server restart service for Metadata server ST1.

```
tanktool> startautorestart ST1
The automatic restart service for server ST1 successfully enabled.
```

Related topics

- "lsautorestart" on page 190
- "Metadata server" on page 8
- "stopautorestart" on page 289

startcluster

Starts all Metadata servers in the cluster and brings them to the full online state.

Syntax

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

Description

Prerequisites:

1. You must have Operator or Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

None of the Metadata servers in the cluster can be running when using this command.

When you start all of the Metadata servers in the cluster using the **startcluster** command, all of the subordinate Metadata servers are brought to the online state. When you start the Metadata server using the **startserver** command, the subordinate Metadata servers in the joining state are placed in their previous state after the master Metadata server is up and running.

Note: A message stating that the cluster started successfully does not necessarily mean that the cluster is online.

Examples

Start the cluster The following example starts all Metadata servers in the cluster.

```
tanktool> startcluster
Cluster started successfully.
```

Related topics

- “quiescecluster” on page 242
- “resumecluster” on page 248
- “startserver” on page 276
- “statcluster” on page 278
- “stopcluster” on page 290

startengine

Starts the CPU and operating system on one or more storage engines.

Syntax

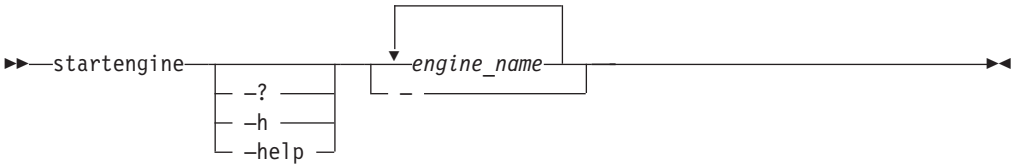

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

engine_name
Specifies the name of one or more storage engines to start.

– Specifies that you want to read the names of one or more storage engines to start from stdin (for example, – << /work/engine_list.txt).

Description

Prerequisite: You must have Operator or Administrator privileges to use the command.

You can stop and restart the local engine, but you cannot start the local engine.

Examples

Start an engine. The following example starts engine ST1.

```
tanktool> startengine ST1
A start request has been sent to Engine ST1.
Tip: Run the lsengine command for current Engine status.
```

Related topics

- “Engines” on page 5
- “lsengine” on page 204
- “restartengine” on page 247
- “statengine” on page 282
- “stopengine” on page 291

startmetadatacheck

Starts the utility that performs a consistency check on the metadata for the entire system or a set of filesets, generates reports in the cluster log, and optionally repairs inconsistencies in the metadata.

Syntax

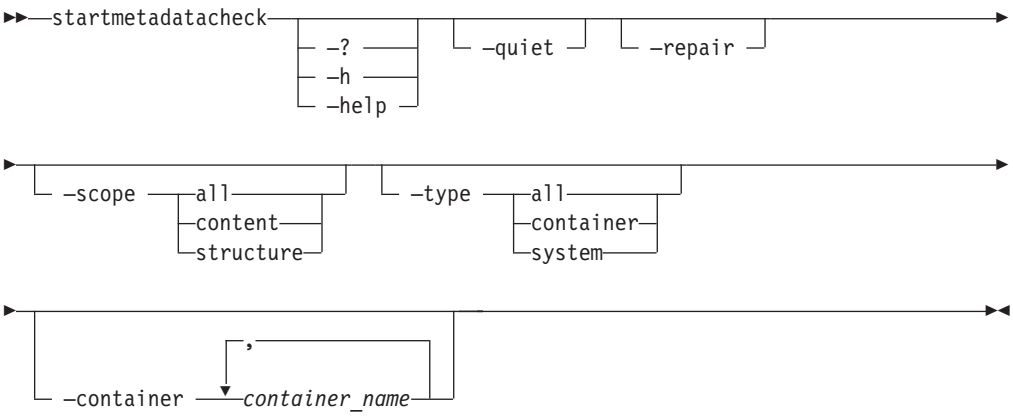

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–quiet
Turns off the prompt to confirm the repair operation.

–repair
Repairs inconsistencies in the metadata and reports the changes in the cluster log. If not specified, repairs are not made.

Note: In some cases, manual intervention might be needed to repair the metadata that requires you to take the cluster offline.

–scope all | content | structure
Specifies the scope of the metadata check. You can specify one of the following values:

all Checks both the metadata content and structure. This is the default value.

content
Checks only the metadata content.

structure
Checks only the metadata structure.

–type all | container | system
Specifies the type of metadata to check. You can specify one of the following values:

all Checks both the system and fileset metadata. This is the default value.

container
Checks only the fileset (or container) metadata.

system
Checks only the system metadata.

– container*container_name*
Specifies the names of one or more filesets (or container) to check, separated by a comma but no spaces. (Do not include a space after the comma.) If not specified, all filesets are checked.

Note: You must specify the **–type container** parameter with this parameter to perform a check on a limited set of filesets. You cannot use this parameter with the **–type all** or **–type system** parameters.

Description

Prerequisites:

1. You must have Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

This command is useful when you suspect inconsistencies in the metadata or after a major restoration of the system.

You must have Administrator privileges to use all parameters in this command. Users with Backup privileges are not allowed to use the –repair parameter.

Note: The metadata checker uses space temporarily in each fileset.

Examples

Check the system-metadata structure The following example checks the system-metadata structure and makes necessary repairs.

```
tanktool> startmetadatacheck –repair –scope structure –type system
Are you sure you want to start a metadata check process? [y/n] y
```

Related topics

- “lscontainer” on page 198
- “stopmetadatacheck” on page 292

startserver

Starts the specified Metadata server.

Syntax

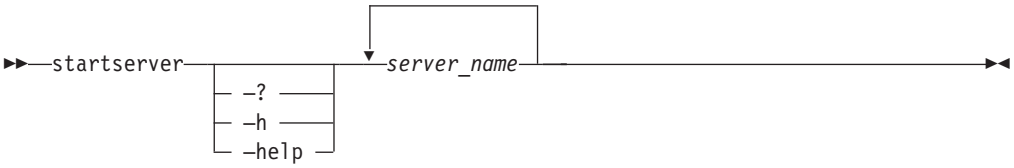

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

server_name
Specifies the name of one or more Metadata servers to start. This Metadata server must be part of the cluster.

Description

Prerequisite: You must have Operator or Administrator privileges to use the command.

Note: If you run this command from an engine hosting a subordinate Metadata server, you can only start the local Metadata server. If you run this command from the engine hosting the master Metadata server, you can start any Metadata server.

If the master Metadata server is not running, you cannot start a subordinate Metadata server.

When you start the Metadata server using the **startserver** command, the subordinate Metadata servers in the joining state are placed in their previous state after the master Metadata server is up and running. When you start the cluster using the **startcluster** command, all of the subordinate Metadata servers are brought to the online state.

The state of the Metadata server matches the state of the cluster. Use the **statcluster** command to check on the state of the cluster, or use the **lsserver** command to check on the state of the Metadata server.

When you start a Metadata server using the **startserver** command, the Metadata server restart service on that Metadata server is enabled only if Metadata server restart service was previously in the STANDBY state. If you manually disable the Metadata server restart service and then stop and restart the Metadata server, the Metadata server restart service ***will not*** become enabled.

Note: A message stating that the Metadata server started successfully does not necessarily mean that the Metadata server is online.

Examples

Start a Metadata servers. The following example starts the Metadata server ST1.

```
tanktool> startserver ST1
Server ST1 started successfully
```

Related topics

- "Metadata server" on page 8
- "lsserver" on page 221
- "statcluster" on page 278
- "statserver" on page 286
- "stopserver" on page 293

statcluster

Displays status, network, workload, and configuration information about the cluster.

Syntax

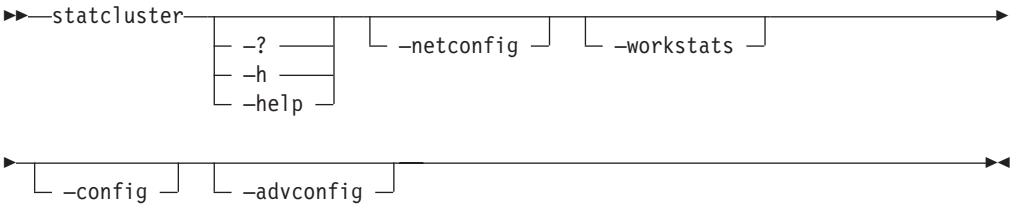

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–netconfig
Displays network configuration information for the cluster.

–workstats
Displays workload statistics related to the total number of transactions, updates, and buffers for the master Metadata server.

–config
Displays cluster configuration settings that, if changed, do not require that you restart the cluster.

–advconfig
Displays cluster configuration settings that, if changed, require that you restart the cluster.

Description

Prerequisite: You must be logged in to the engine hosting the master Metadata server to run this command.

If you do not specify a parameter, this command displays the following default information:

- Name of the cluster.
- ID of the cluster.
- Current state of the cluster. Possible states are:

  Forming
  The cluster has a master Metadata server and is in the process of forming. The first time a cluster is formed, its initial state is FORMING.

  Fully quiescent
  One or more Metadata servers in the cluster are in the FULLY QUIESCENT state.

  Not running
  One or more Metadata servers in the cluster are not added.

Offline
One or more Metadata servers in the cluster are offline.

Online
All Metadata servers in the cluster are online.

Partly quiescent
One or more Metadata servers in the cluster are in the PARTLY QUIESCENT state.

Unknown
The state of the cluster is not known.

- Target state to which the cluster is switching.
- Timestamp of the last current state change.
- Timestamp of the last state change.
- Number of Metadata servers in the cluster.
- Number of Metadata servers that are actively participating in the cluster (in the offline, online, or quiescent state).
- Committed software version.
- Software version.
- Timestamp when the latest software was committed.
- Software commit status (In Progress or Not In Progress).
- Timestamp of the last installation.

If you specify the **–netconfig** parameter, this command displays the following information in addition to the default information:

- Name of the master Metadata server.
- IP address.
- Cluster port number.
- Heartbeat port number.
- Client-Metadata server port number.
- Administrative port number.

If you specify the **–workstats** parameter, this command displays the following information in addition to the default information:

- Number of system-update transactions involving the creation, modification, and deletion of system objects.
- Number of system-object transactions, involving the creation, modification, and deletion of objects. The system objects include storage pools, filesets, volumes, policies, and engines.
- Current number of buffers, which are used for system metadata activity, that contain data but are available for reuse (clean buffers).
- Current number of buffers, which are used for system metadata activity, that contain data awaiting I/O to disk (dirty buffers).
- Current number of buffers, which are used for system metadata activity, that are not in use (free buffers).
- Current total number of buffers for system metadata activity.

If you specify the **–config** parameter, this command displays the following information in addition to the default information:

- Pool-space reclamation interval, in minutes.

- List of clients with root or administrative privileges.
- Size, in 4-KB pages, of the master Metadata server buffer.
- Size, in pages, of the subordinate Metadata server buffer.
- Metadata server workload-process limit.
- Administrative process limit.

If you specify the **–advconfig** parameter, this command displays the following information in addition to the default information:

- Threshold that specifies the maximum number of missed network heartbeats.
- Maximum number of missed network heartbeats.
- Network heartbeat interval, in milliseconds.
- Cluster timeout, in milliseconds.
- Maximum number of retries to client.
- Client timeout, in milliseconds.
- Client lease period, in seconds.
- Client-lease-period multiplier.
- Metadata server timeout, in milliseconds.

Examples

List network information for the cluster The following example displays the default and network information for the cluster:

```
tanktool> statcluster -netconfig
Name                          sanfs
ID                            2802
State                         ONLINE
Target State                  FULLY QUIESCENT
Last State Change             Sep 24, 2003 3:31:52 PM
Last Target State Change      Sep 24, 2003 2:40:20 PM
Servers                       4
Active Servers                3
Committed Software Version    1.1.4.2
Software Version              1.00
Last Software Commit          Jan 3, 2003 1:40:20 PM
Software Commit Status        In Progress
Installation Date             Jan 3, 2003 4:56:59 PM

===============Network Configuration===============
Master Server                 ST1
IP                            128.0.0.1
Cluster Port                  1737
Heartbeat Port                1738
Client-Server Port            1700
Admin Port                    1800
```

Related topics

- “chclusterconfig” on page 172
- “quiescecluster” on page 242
- “resumecluster” on page 248
- “startcluster” on page 272
- “stopcluster” on page 290

statcontainer

Displays the number of started and completed transactions for the filesets (containers) being served by the local Metadata server.

Syntax

Description

Note: This command is run from the shell prompt. It is not run inside of tanktool.

A *transaction* is a work request that is handled by the Metadata server. You would use the number of transactions performed by each fileset on the Metadata server to aid in balancing the workload among all of the Metadata servers in the cluster.

This command displays the following information for each fileset:
- Name of the fileset.
- Number of transactions that have been initiated since the Metadata server started (including those transactions that are initiated during startup). This count includes the transactions that are in progress, stopped, and completed.
- Number of transactions that completed successfully.

Examples

Displaying number of transactions The following example displays the number of started and completed transactions for filesets being served by the local Metadata server.

```
tanktool> statcontainer

Name        Transactions Started   Transactions Completed
=========================================================
cnt_A                       12765                    12751
cnt_B                        9478                     9465
```

Related topics

- "attachcontainer" on page 166
- "chcontainer" on page 174
- "detachcontainer" on page 183
- "lscontainer" on page 198
- "mkcontainer" on page 230
- "rmcontainer" on page 251

statengine

Displays status information about a specific storage engine.

Syntax

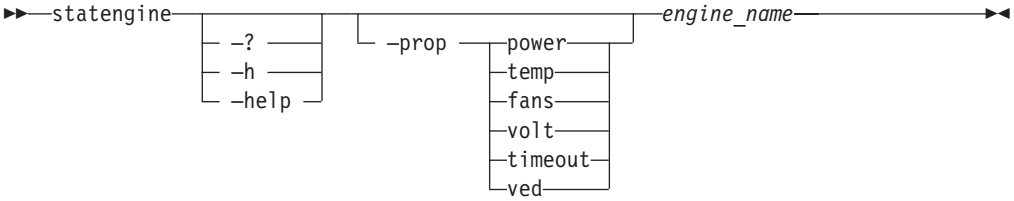

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–prop
Specifies the property status to display. Possible values are:

power Displays power status for the specified engine and operating system.

temp Displays temperature status for the components of the specified engine.

fans Displays fan status for the specified engine.

volt Displays voltage status for the specified engine.

timeouts
Displays timeout information for the certain hardware and software functions in specified engine environment.

ved Displays vital engine data for the specified engine.

engine_name
Specifies the name of the storage engine to display.

Description

If you do not specify a property parameter, this command displays the following default information for each engine:

- Engine IP address.
- Engine name.
- Boot state. Possible values are:

  Before POST
  The engine is powered on but has not started the power-on self test (POST).

  Booted Flash
  The engine has read the System BIOS but has not started loading the operating system.

  Booting OS
  The engine has stated but not completed loading the operating system.

  CPUs Held in Reset
  The engine has been reset after a hardware fault.

In POST
The engine is running the POST.

In OS The engine is running in the normal state.

Stopped in POST (error detected)
The engine is powered on but has not completed the POST due to an error.

Unknown / Power Off
The boot state is unknown. The engine could be powered off.

- Temperature state. Possible values are:

  – The temperature thresholds are not set.

  Normal
  The temperatures of all components in all engines are below the Warning threshold.

  Warning
  The temperature of one or more engine components are above the Warning threshold.

  Error The temperature of one or more engine components is above than the Soft Shutdown threshold.

  Unknown
  The RSA card could not be accessed.

- Fan state. Possible values are:

  Normal
  All fans are operating above 15% of its fan-speed capacity.

  Warning
  One or more fans are operating below 15% of its fan-speed capacity.

  Error The RSA card could not be accessed.

- Voltage state. Possible values are:

  – The voltage thresholds are not set.

  Normal
  The voltages of all components in all engines are above the Warning Low Voltage threshold and below the Warning High Voltage threshold.

  Warning
  The voltage of one or more engine components is below the Warning Low Voltage threshold or above the Warning High Voltage threshold.

  Error The RSA card could not be accessed.

If you specify the **–prop power** property parameter, this command displays the following for each engine:

- Engine IP address.
- Engine name.
- Power state (On or Off).
- Number of hours the engine has been powered on.
- Number of times the engine has been restarted. This counter is cleared when the Advanced System Management (ASM) device is cleared to factory defaults.
- Current time on the ASM device's local clock. This time is independent of the time on the engine and is the time used to schedule a power off.

If you specify the **–prop temp** property parameter, this command displays the following for each engine component:

- Engine component. This command displays information for the following components. (Note that the actual component labels and numbers may vary depending on your environment.):

  CPU1

  CPU2

  DASD1

  DASD2

  Ambient (overall temperature)
- Temperature state per component.
- Current temperature, in degrees Celsius.
- Warning threshold, in degrees Celsius.
- Soft Shutdown threshold, in degrees Celsius.
- Hard Shutdown threshold, in degrees Celsius.

If you specify the **–prop fans** property parameter, this command displays the following for each engine:

- Fan name. This command displays information for each of the eight fans, labeled Fan 1 through Fan 8.
- Fan state.
- Fan speed as a percentage of the maximum capacity speed.

If you specify the **–prop volt** property parameter, this command displays the following for each engine:

- Engine component. This command displays information for the following components. (Note that the actual component labels and numbers may vary depending on your environment.):

  VRM 1 (1.5 V)

  VRM 2 (1.5 V)

  System board (12 V)

  System board (5 V)

  System board (3.3 V)

  System board (2.5 V)

  System board (1.5 V)
- Voltage state.
- Voltage.
- Description.
- Warning Low voltage threshold.
- Warning High voltage threshold.

If you specify the **–prop timeouts** property parameter, this command displays the following for each engine:

- POST Watchdog threshold. This is the amount of time the engine will wait for POST to complete before sending an alert and automatically restarting the system. If the threshold is not set, the value is –.
- Loader Watchdog threshold. This is amount of time the engine will wait for the operating system to load before sending an alert and automatically restarting the system. If the threshold is not set, the value is –.

- OS Watchdog timeout. This is the amount of time for the operating system to respond before sending an alert and automatically restarting the system. If the threshold is not set, the value is –.
- OS Watchdog check interval. This is the amount of time between the OS Watchdog checks. If the threshold is not set, the value is –.
- Power Off Delay threshold. The amount of time to wait for the operating system to shut down before powering off the system. If the threshold is disables, the value is –.

If you specify the **–prop ved** property parameter, this command displays the following for each engine:

- Engine model number.
- Engine serial number.
- Unique universal identifier, which uniquely identifies the engine.
- Firmware revision. This is the revision numbers of the application firmware and startup ROM firmware.
- Logical device's firmware revision date.
- File names of the application and startup ROM firmware.
- Build identifier of the application and startup ROM firmware.

Examples

Display power status for an engine The following example displays temperature status information for the engine ST1.

```
tanktool> statengine –prop power ST1

Engine IP            39.47.25.19
Name                 ST1
Power State          Off
Power-On Hours       256 hours
Restarts             56
Current ASM Time     Aug 12, 2003 4:59:47 PM
```

Related topics

- "Engines" on page 5
- "startengine" on page 273
- "lsengine" on page 204
- "stopengine" on page 291

statserver

Displays status, configuration, and workload information for a specific Metadata server in the cluster, if issued from the master Metadata server. Lists status, configuration, and workload information for the local Metadata server, if issued from a subordinate Metadata server.

Syntax

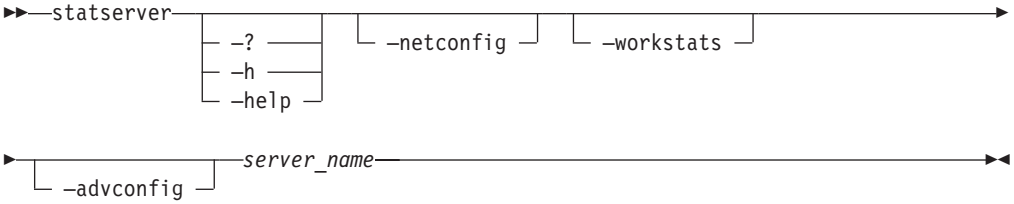

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–netconfig
Displays network configuration information for the specified Metadata server.

–workstats
Displays statistical information for the local Metadata server.

–advconfig
Displays advanced settings, which require a reboot when changed, for the specified Metadata server.

server_name
Specifies the Metadata server to display. If you issue this command from a subordinate Metadata server, this must be the name of the local Metadata server.

Description

Note:

> If you run this command from an engine hosting a subordinate Metadata server, you can display information about only the local Metadata server. If you run this command from the engine hosting the master Metadata server, you can display information about any Metadata server.

If you do not specify any parameters, the following default statistics are displayed:
- Metadata server name.
- Role of the Metadata server (Master or Subordinate).
- Pending software version.

If you specify the **–netconfig** parameter, the following statistics are displayed in addition to the default statistics:
- Master Metadata server name.
- Metadata server-to-Metadata server transport protocol (UDP)
- Client-to-Metadata server transport protocol (TCP)
- IP address.

- Cluster port number.
- Heartbeat port number.
- Client-server port number.
- Administrative port number.

If you specify the **–workstats** parameter, the following statistics are displayed in addition to the default statistics:

- Number of file-system-update transactions, involving creating, modifying, and deleting system objects.
- Number of file-system transactions, involving reading, creating, modifying, and deleting objects. The system objects include storage pools, containers, volumes, policies, and engines.
- Number of buffers that contain data awaiting I/O to disk (dirty buffers).
- Number of buffers that contain data but are available for reuse (clean buffers).
- Number of buffers that are not in use (free buffers).
- Total number of available buffers.
- Number of session locks held by the engine hosting the Metadata server. Clients acquire session locks to perform file management operations.
- Number of data locks held by the engine hosting the Metadata server. Clients hold data locks to cache data pages and attributes of files and to cache read-only attributes and contents of directories and links.
- Number of byte-range locks held by the engine hosting the Metadata server. Clients use byte-range locks to implement POSIX, SYSV and Berkeley lock system calls.

If you specify the **–advconfig** parameter, the following statistics are displayed in addition to the default statistics:

- File-space reclamation process limit.
- Synchronous commit mode
  sync Schedules and buffers metadata changes instead of an immediate write.
  async Metadata changes are written to disk immediately.
- List of current profile parameters being used.
- List of current trace parameters being used as output.
- Time interval, in seconds, at which the user and system-time statistics are gathered.

Examples

List information about the Metadata server The following example displays all information about the server:

```
tanktool> statserver –netconfig ST1

Server                                ST1
Server Role                           Subordinate
Most Current Software Version         V1.003
==================Network Configuration==================
Master Server                         ST4
Server-to-Server Transport Protocol   UDP
Client-to-Server Transport Protocol   TCP
IP                                    128.0.0.1
Cluster Port                          1737
Heartbeat Port                        1738
Client-Server Port                    1700
Admin Port                            1800
```

Related topics

- "Metadata server" on page 8
- "lsserver" on page 221

stopautorestart

Disables the Metadata server from restart automatically if it is down.

Syntax

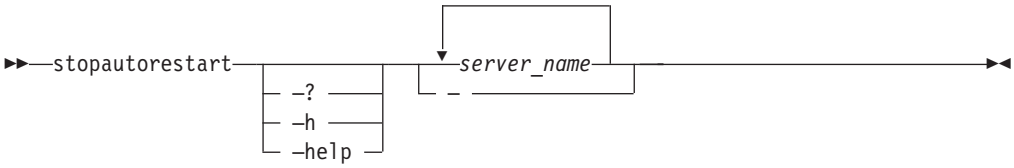

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

server_name
Specifies the names of one or more Metadata servers to disable from restarting automatically.

– Specifies that you want to read the names of one or more Metadata servers to enable from restarting automatically from stdin (for example, – << /work/server_list.txt).

Description

Prerequisite:: You must have Administrator privileges to use the command.

Note: If you run this command from an engine hosting a subordinate Metadata server, you can stop the Metadata server restart service on only the local Metadata server. If you run this command from the engine hosting the master Metadata server, you can stop the Metadata server restart service on any Metadata server.

Examples

Disable the automatic-restart service The following example disables the automatic-restart service for Metadata server ST1.

```
tanktool> stopautorestart ST1
The automatic restart service for server ST1 successfully disabled.
```

Related topics

- “lsautorestart” on page 190
- “Metadata server” on page 8
- “startautorestart” on page 270

stopcluster

Stops all Metadata servers in the cluster gracefully.

Syntax

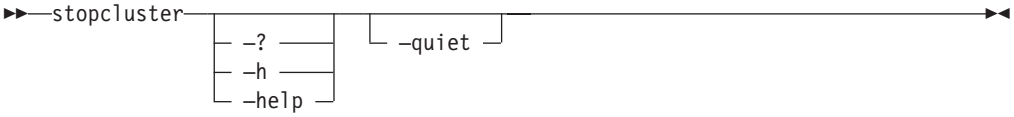

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–quiet
Turns off the prompt to confirm the change operation

Description

Prerequisites:

1. You must have Operator or Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

This command does not shut down the engine on which the specified Metadata server runs.

When you stop the master Metadata server using the **stopserver** command, the subordinate Metadata servers are abruptly moved to the joining state. When you stop all of the Metadata servers in the cluster using the **stopcluster** command, all of Metadata servers are brought down gracefully.

Examples

Stop the cluster The following example stops the cluster gracefully.

```
tanktool> stopcluster
Are you sure you want to shut down the cluster? [y/n] y
Cluster shut down gracefully.
```

Related topics

- "quiescecluster" on page 242
- "resumecluster" on page 248
- "startcluster" on page 272
- "statcluster" on page 278

stopengine

Shuts down the operating system and powers off one or more storage engines.

Syntax

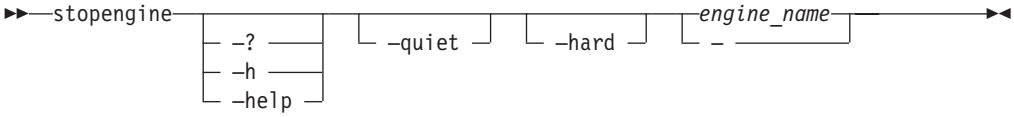

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–hard
Powers off the specified engines without first shutting down the operating system. If not specified, this command shuts down the operating system before powering off the storage engines.

–quiet
Turns off confirmation prompts for this command.

engine_name
Specifies the name of the storage engines to power off.

– Specifies that you want to read the names of one or more storage engines to power off from stdin (for example, – << /work/engine_list.txt).

Description

Prerequisite: You must have Operator or Administrator privileges to use the command.

You can stop and restart the local engine, but you cannot start the local engine.

Examples

Power off an engine. The following example shuts down the operating system and then powers off engine ST1.

```
tanktool> stopengine ST1
Are you sure you want to stop engine ST1? [y/n] y
A start request has been sent to Engine ST1.
Tip: Run the lsengine command for current Engine status.
```

Related topics

- “Engines” on page 5
- “lsengine” on page 204
- “restartengine” on page 247
- “startengine” on page 273
- “statengine” on page 282

stopmetadatacheck

Stops the metadata check utility that is currently in progress.

Syntax

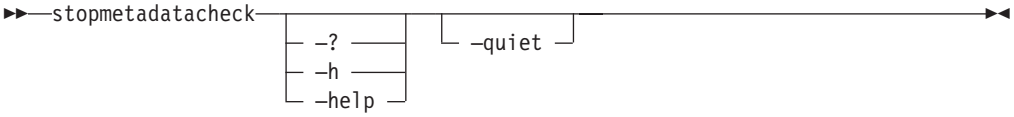

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–quiet
Turns off the prompt to confirm the stop operation.

Description

Prerequisites:

1. You must have Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

Examples

Stop the metadata check process The following example stops the current metadata check process.

```
tanktool> stopmetadatacheck
Are you sure you want to stop the metadata check process? y/n  y
The metadata check is stopping.  This may take a few minutes to complete.
```

Related topics

- "lscontainer" on page 198
- "startmetadatacheck" on page 274

stopserver

Shuts down a subordinate Metadata server gracefully.

Prerequisites:

1. You must have Operator or Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

Syntax

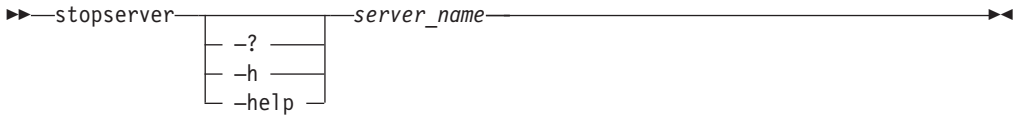

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

server_name
Specifies the name of the subordinate Metadata server to shut down. This Metadata server must be part of the cluster.

Description

To use this command, you must be logged into the master Metadata server and the master Metadata server must be running.

This command does not shut down the engine on which the specified Metadata server runs and does not remove the Metadata server from the cluster.

When you stop the master Metadata server using the **stopserver** command, the subordinate Metadata servers are abruptly moved to the JOINING state. When you stop all of the Metadata server in the cluster using the **stopcluster** command, all of Metadata servers are brought down gracefully.

When you stop a Metadata server using the **stopserver** command, the Metadata server restart service on that Metadata server is changed to the STANDBY state only if Metadata server restart service was previously in the ON state.

Examples

Stop a Metadata server The following example stops the Metadata server ST1.

```
tanktool> stopserver ST1

Server ST1 shutdown successfully.
```

Related topics

- “Metadata server” on page 8
- “lsserver” on page 221
- “startserver” on page 276
- “statserver” on page 286
- “stopcluster” on page 290
- “stopengine” on page 291

suspendvol

Suspends one or more volumes so that the Metadata server cannot allocate new data on the volumes.

Syntax

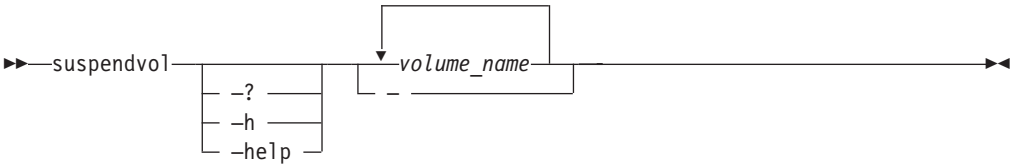

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

volume_name
Specifies the names of one or more volumes to suspend.

– Specifies that you want to read the names of one or more volumes to suspend from stdin (for example, – << /work/vol_list.txt).

Description

Prerequisites:

1. You must have Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

Examples

Suspend a volume The following example suspends activity on volume vol2.

```
tanktool> suspendvol vol2
Volume vol2 suspended.
```

Related topics

- “mkvol” on page 240
- “chvol” on page 178
- “lsvol” on page 227
- “activatevol” on page 163
- “reportvolfiles” on page 245

upgradecluster

Verifies that all the engines in the cluster have been upgraded to the new version of the software and then initiates the cluster upgrade process.

Syntax

Parameters

–quiet
Turns off the prompt to confirm the upgrade operation.

Description

Prerequisites:

1. You must have Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

This command only works after all engines have been upgraded to the same software version.

You do not need to stop the cluster to perform the upgrade.

Examples

Upgrade the cluster The following example upgrades the cluster:

```
tanktool> upgradecluster
Are you sure you want to upgrade the cluster software? [y/n] Y
Cluster upgrade successful.
```

Related topics

- “startcluster” on page 272
- “quiescecluster” on page 242
- “resumecluster” on page 248
- “statcluster” on page 278

usepolicy

Directs the Metadata server to make an existing policy the active policy and applies its rules to all subsequent file creations.

Syntax

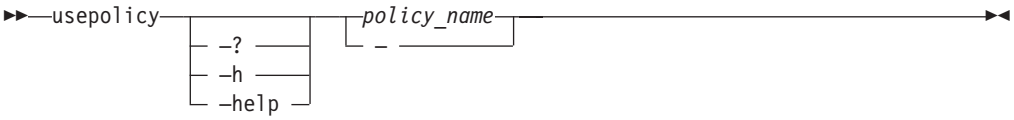

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

policy_name
Specifies the name of the policy to activate.

– Specifies that you want to read the name of the policy to activate from stdin (for example, – << /work/policies_list.txt).

Description

Prerequisites:

1. You must have Administrator privileges to use the command.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

Examples

Activate a policy The following example activates the policy *test_policy*.

```
tanktool> usepolicy  test_policy
Policy test_policy is now the active policy.
```

Related topics

- "catpolicy" on page 171
- "lspolicy" on page 213
- "mkpolicy" on page 236
- "rmpolicy" on page 255

Client commands

There is a command-line interface for each client operating system that SAN File System supports.

Related topics:

- "AIX-client commands"
- "Windows-client command" on page 315

AIX-client commands

The following table provides a brief description for each AIX-client command.

Note: You must have root privileges to use these commands.

Command	Description
Migration	
migratedata	Migrates data to SAN File System.
Status	
stfsstatus	Displays the version of the file-system drive for the specified virtual client.
Volumes and LUNs	
stfsdisk	Scans the SAN File System for new and removed volumes.
Virtual client setup and removal	
rmstclient	Unmounts the global namespace, removes the SAN File System client, and unloads the file-system driver.
setupstclient	Loads the file-system driver from a standard location, creates the SAN File System client, and mounts the global namespace.
stfsclient	Creates or destroys a virtual client.
stfsdriver	Loads the file-system driver as a kernel extension.
stfsmount	Mounts the global namespace.
stfsumount	Unmounts the global namespace.

migratedata

Migrates data to SAN File System.

Syntax:

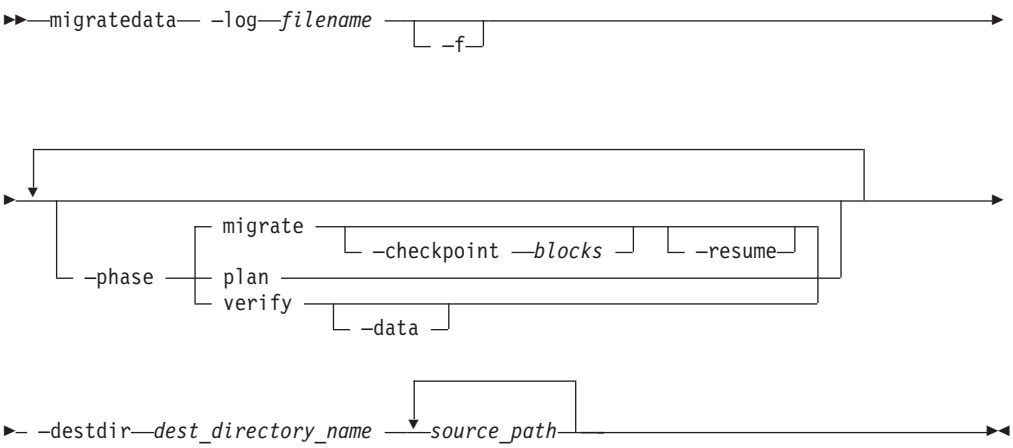

Parameters:

–log *filename*
Specifies the location of a file in which to log migration activities, warnings, and errors. When used with the **–plan migrate –resume** parameter, the **–log** parameter specifies the log file from which to read information about the last completed block or file.

Attention: You must specify the correct log file with migrate –resume and verify that the source and destination directories specified on the command line match those in the log file.

- If you specify an incorrect log file and the **–f** parameter, –resume displays a warning and overwrites the target file-system data with wrong information.
- If you specify an incorrect log file, but do not specify the **–f** parameter, this command displays an error and exits.

–f Specifies that the migration should continue even if there is an error with a file. If specified with the **–phase migrate** parameter, this command skips any files with errors, but continues with the migration process. If not specified, an error results in the entire migration being stopped before the file that caused the error. You can then restart the migration after fixing the error.

If specified with the **–phase verify** parameter, this command adjusts any missing metadata attributes, such as permissions and times. If there is a mismatch in size, however, this command will not remigrate the file.

–phase
Specifies the migration phase to run. Choices include:

plan Gathers information about the available system resources (available memory, number of CPUs, size of the source tree and space available on the destination file system), copies sample files from source directory to estimate transfer rates, and provides an estimated time for the migration of the data set.

migrate
Reads data from source file system and writes the data to the destination file system. Although not required, for large data sets, you

should run this command in planning mode first. You can stop the migration process at any point and resume from the last completed file or block (using the **–resume** parameter).

This is the default value.

verify Verifies the integrity of the migrated data using the Message Digest 5 verification algorithm on the contents of the file, as well as verifying consistency of the metadata (such as owner and modification time stamp settings) between the source and destination files.

You can specify more than one phase. For example, to plan, migrate, and verify the data, specify **–phase plan –phase migrate –phase verify**. Although you can specify the phases in any order, this command always estimates the completion time, migrates data, and then verifies the migrated data.

If the **–phase** parameter is not specified, this command runs only the migration phase.

–checkpoint *blocks*
Writes a checkpoint in the log file after each specified number of blocks of file data has completed the migrate phase. (The block size depends on the client platform.) For example, if you specify **–checkpoint 20**, this command makes an entry in the log file each time 20 blocks of file data is migrated. On a platform with a block size of 16 MB, this command writes to the log file after each 2 560-MB of the file data has been migrated. If the process is interrupted, you can resume the migration at the place it left off. If unspecified, the **migratedata** command makes an entry in the log file after each complete file has been migrated. You can resume the migration at the point of the last migrated file.

–resume
Resumes the migration from the last completed block or file (logged in the log file specified by –log). If the log file indicates that some files in the source directory are migrated and this parameter is not specified, this command restarts the migration process from the beginning (performs a fresh migration).

–data
Verifies every block of source data (file data and metadata) with the destination data. If not specified, this command verifies only the metadata unless there is a mismatch in the file attributes, in which case this command then verifies the file data.

Note: Verifying all data is very time consuming and can take as long as the migration itself.

–destdir *dest_directory_name*
Specifies the name of the destination directory for the migrated data. The directory can either exist or be a new directory name. IBM recommends that you create the directory before beginning the migration process. If the directory does not exist, this command creates the directory using the default permissions.

source_path
Specifies one or more paths of directories or files to migrate.

Description:

Prerequisite:

1. You must have root privileges on AIX® or Administrative privileges on Windows to use this command.

2. All storage pools, all filesets, and at least one policy must be set up. All activity (from applications, such as database servers and application servers, or users) that modifies data on the source and destination file systems must be stopped and remain stopped to guarantee consistency of the migrated data.
3. The destination directory must exist with correct set of permissions and appropriate storage policies must be configured.

Examples: Migrating data from a client for AIX This example migrates data from the work/capital directory on the client machine to the sanfs/cnt1 directory in the global namespace. A checkpoint is written to the mgrt_capital.log log file each time 20 blocks of file data is migrated.

```
migratedata –log /mgrtlogs/mgrt_capital.log –phase migrate –checkpoint 20
–destdir /mnt/tank/sanfs/cnt1 work/capital
```

Migrating data from a client for Windows This example migrates data from the C:\Capital directory on the client machine to the sanfs\cnt1 directory in the global namespace. A checkpoint is written to the mgrt_capital.log log file each time 20 blocks of file data is migrated.

```
migratedata –log c:\mgrtlogs\mgrt_capital.log –phase migrate –checkpoint 20
–destdir t:\cnt1 C:\Capital
```

rmstclient

Unmounts the global namespace, removes the virtual client for AIX, and unloads the file system driver from the local client machine.

Syntax:

Parameters:

–prompt
Prompts for required parameters, using values from the configuration file, if available.

–noprompt
Runs silently, using parameters from the configuration file (/usr/tank/client/config/stclient.conf). If a required parameter is not available, the command exists with an error.

Description:

Prerequisite:

1. You must have root privileges to use this command.
2. You must unmount the SAN File System before invoking this command.

Examples: Remove a client for AIX The following example removes the local SAN File System client for AIX without prompting.

```
rmstclient –noprompt
```

Related topics:

- “setupstclient” on page 302
- “stfsmount” on page 311
- “stfsumount” on page 314
- “Clients” on page 10
- “UNIX-based clients” on page 12

setupstclient

Sets the client configuration and saves the configuration, and optionally loads the file-system driver, creates the client for AIX, and mounts the global namespace.

Syntax:

Parameters:

–prompt
Prompts for required parameters, using values from the configuration file, if available.

–noprompt
Runs silently, using parameters from the configuration file (/usr/tank/client/config/stclient.conf). If a required parameter is not available, the command exists with an error.

Description:

Prerequisite: You must have root privileges to use this command.

This command prompts you for information necessary to set the client configuration or retrieves the information from the configuration file.

If you do not specify a parameter, this command run silently using values from the configuration files as defaults and prompts for any required information.

Examples: Setup a client for AIX The following example loads the file-system driver, creates the client, and mounts the global namespace.

```
setupstclient
```

Related topics:

- "rmstclient" on page 301
- "Clients" on page 10
- "UNIX-based clients" on page 12

stfsclient

Creates or destroys a virtual client.

Syntax:

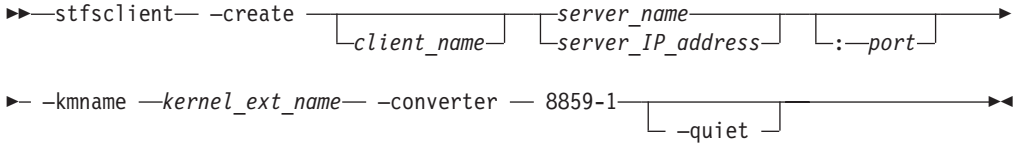

or

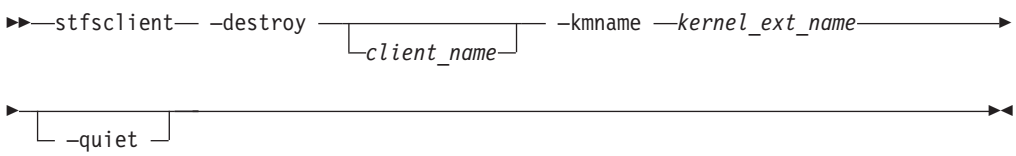

Parameters:

–create
Creates a new virtual client.

–destroy
Destroys an existing virtual client.

client_name
Identifies the unique name of the virtual client that you want to create or destroy. The default client name is the host name of the client system.

server_name
Specifies the host name of a Metadata server in the SAN File System. The Metadata server that you specify informs the global namespace image of all other Metadata servers.

This parameter is not required if this is not the first mount for a particular virtual client.

server_IP_address
Specifies the IP address, in dotted decimal notation, of a Metadata server in the SAN File System.

port
Specifies the port number of the specified Metadata server. The default is 1700.

–kmname *kernel_ext_name*
Identifies kernel-extension name of the file-system-driver instance associated with the virtual client.

The file-system driver is loaded as a kernel extension. To identify the instance of the file-system driver, you identify the kernel extension. The kernel-extension name is the same as name and location of the file-system driver that was used to load the driver (for example, /usr/tank/client/bin/stfs for AIX).

–devices
Determines which devices (also called disks or LUNs) that the virtual client

considers as SAN File System volumes. The default is the value of the STFS_DEVICES environment variable or, if that is not set, ″–devices=pat=/dev/rhdisk*.

In addition to creating the virtual client, this command discovers which disks, or candidates, are available to the virtual client as volumes and transmits the candidate list to the virtual client. The **–devices** parameter controls the candidates list.

dir=*directory*
The candidates list is made up of those devices that have device special files in the specified directory (for example: –devices=dir=/dev/stfsdisk).

The easiest way to mount the global namespace is to specify –devices=pat=/dev/rhdisk*, which looks at every SCSI-disk-like device in the system and whatever looks like a SAN File System disk is accessed when the Metadata server refers to that disk's SAN File System disk identifier.

If you want the client to be more selective about what disks it considers available, you can create a /dev/stfsdisk directory, put device-special files (or symbolic links) for your candidates in it, and just let –devices=dir=/dev/stfsdisk default.

pat=*pattern*
The candidates list is made up of those devices that have device-special files whose file specifications match the specified pattern. You can use * wildcards in the last (filename) component but not in the directory components (for example, –devices=pat=/dev/rhdisk*).

none The candidates list is empty. Use this value when you want to establish the candidate list with a separate command, perhaps using a selection method more sophisticated than the stfsclient command offers.

–quiet
Specifies not to display informational messages. This parameter does not affect error messages.

Description:

Prerequisite: You must have root privileges to use this command.

This command creates or destroys a virtual client. A *virtual client* is an entity that communicates with a Metadata server and, indirectly, with other SAN File System client. In this release, only one virtual client is supported per client machine. The terms virtual client and client can be used interchangeably.

A virtual client is associated with exactly one SAN File System. There is one file cache and one set of disk candidates per virtual client. Each virtual client running on the same system is as separate as if it were running on a different system. They share nothing except the file-system drive code that they execute.

A SAN File System virtual client is uniquely identified in the context of its file-system driver, and in the context of its SAN File System, by its client name.

To use the files in a global namespace, the virtual client must have a global namespace image. Creating a global namespace image makes the directory

structure in the global namespace appear in the client's file structure. To create a global namespace image, use the **stfsmount** command.

A client can access and create data that is stored in a global namespace. Each virtual client can access data on multiple images in the same global namespace.

The client considers a file to be one file even if it appears with two different file names in two different global namespace images.

Examples: Create a virtual client The following example creates a virtual client.

```
stfsclient –create MDS1:1700 –kmname /usr/tank/client/bin/stfs
–converter 8859-1
```

Related topics:

- "Clients" on page 10
- "stfsmount" on page 311
- "stfsdriver" on page 309
- "stfsdisk" on page 306
- "UNIX-based clients" on page 12

stfsdisk

Controls the SAN File System volumes (disks) that a client can access.

Syntax:

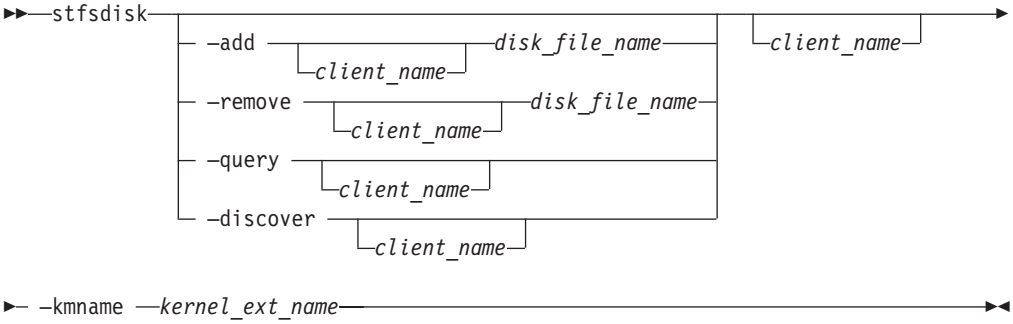

Parameters:

–add
Adds the specified disk-specific file name to the disk-candidate list. If the disk-specific file name is already in the list, the client performs a disk discovery procedure on it, updating its database if indicated.

–remove
Removes the specified disk-specific file name from the disk candidate list. If the disk-specific file name is not in the list, this command does nothing, but does not consider it an error.

–query
Displays the list of disk-specific file names in the current disk-candidate list and the status of each. Possible status values are:

ACTIVE
Indicates that the disk is a valid SAN File System user-data volume and is available to the client to perform file reads and write operations.

INACTIVE
Indicates that the disk is not a valid SAN File System user-data volume and is not available to the client to perform file reads and writes. A disk can be inactive if there were I/O errors when the client tried to access the disk, if the disk does not contain a SAN File System label, or if the disk's SAN File System label says it is something other than a user-data volume.

–discover
Rebuild the database of usable disks by going through the current candidate-disk list and attempting to access each disk, determine if it is a valid SAN File System user-data volume, and read its SAN File System global disk ID. If a disk has become accessible or inaccessible, or changed its identity since the last time this disk-discovery procedure was run, the virtual client updates its candidate-disk list accordingly.

This parameter causes the disk-discovery procedure to start. The procedure typically ends before the disk-discovery procedure completes. While the disk-discovery procedures are in progress, any file-system access that would fail because the virtual client cannot find a specific disk will wait until the disk-discovery procedure completes, and then proceed on the basis of the new disk-accessibility information.

disk_file_name
Specifies the file name of the disk to add to or remove from the disk candidate list. This must be a raw disk file such as /dev/vpath0, not a logical volume file such as /dev/hdisk0.

client_name
Specifies the name of the virtual client whose disk-access you are controlling.

–kmname *kernel_ext_name*
Identifies kernel-extension name of the file-system-driver instance associated with the client.

The file-system driver is loaded as a kernel extension. To identify the instance of the file-system drive, you identify the kernel extension. The kernel-extension name is the same as the name and location of the file-system driver that was used to load the driver (for example, /usr/tank/client/bin/stfs for AIX). Note that the kernel extension name might not be unique.

Description:

Prerequisite: You must have root privileges to use this command.

A client reads and writes files by accessing the disks on which the file data resides. To control which disks that a client can access, SAN File System identifies that disk by a SAN File System global disk identifier, and the disk-access subsystem associates that identifier with the name that the AIX operating system uses to identify that disk. The disk-access subsystem maintains a database that correlates global-disk identifiers with AIX device numbers. When the client needs to access a data block of a file, it consults that database.

The disk-access subsystem maintains the database by reading certain disks at certain times and looking for a SAN File System global disk identifier. If it finds the identifier, it determines whether the disk is a SAN File System user-data volume. If the disk is a volume, it adds the disk to its database.

The set of disks that the disk-access subsystem searches is called the *disk-candidate list*. The **stfsclient** command creates the disk-candidate list when it creates the virtual client. You can modify the list using the **–add** and **–remove** parameters.

The candidate-disk list is a list of unique disk-special file names. Because a disk can be referred to by more than one disk-special file name, the list is not strictly a list of unique devices. Actually examining disks and updating the database of valid user-data volumes is separate from maintaining the candidate-disk list.

When you add a disk to the candidate-disk list, the client immediately tries to read it and adds it to the database. But the disk becomes and stays a candidate regardless of the results of that operation.

You can force the client to rescan the entire list of candidate disks using the **–discover** parameter. The client updates its database of user-data volumes according to the results of this discovery, adding and removing disks as necessary. The results of the discovery do not affect the candidate-disk list, however.

Note that device file names can change as the client runs. Such a change has no effect on the client unless something causes a disk-discovery procedure to run. For example, if you add /dev/rhdisk35 as a candidate disk, and the client successfully identifies it as a SAN File System user-data volume, and then you delete /dev/rhdisk35, the client continues accessing that disk as before. The disk

/dev/rhdisk35 continues to be a candidate. But the next time a disk-discovery procedure runs, the candidate will be found invalid and the client will no longer have access the disk.

Examples: Query the disk-candidate list The following example queries disk-candidate list and displays the status of each disk.

```
stfsdisk –query –kmname /usr/tank/client/bin/stfs
```

Related topics:

- "Clients" on page 10
- "UNIX-based clients" on page 12
- "stfsclient" on page 303
- "stfsdriver" on page 309

stfsdriver

Loads the file-system driver as a kernel extension.

Syntax:

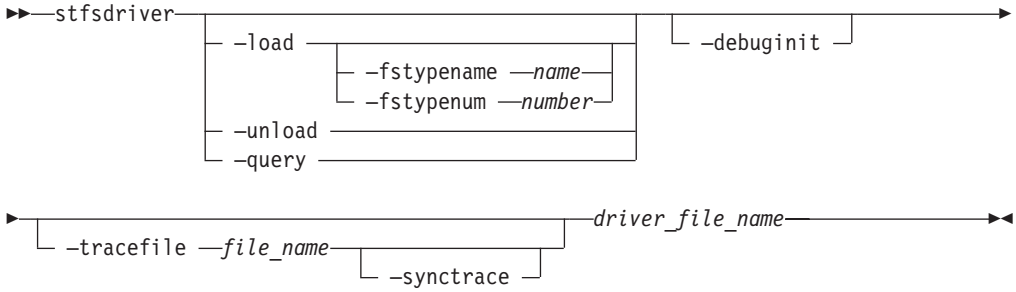

Parameters:

–load
Loads the kernel extension and create an instance of the file-system driver.

–unload
Unloads the kernel extension and destroys the instance of the file-system driver.

–query
Displays information about the kernel extension matching the specified criteria. For example, you can query the kernel extension ID to use in commands instead of the kernel module name.

–fstypename *name*
Specifies the name of the file-system type to use for the file-system-driver instance. This name relates to a specific file-system-type number. The file /etc/vfs maps the file-system-type name to the number.

If you do not specify a file-system-type name or number, the system defaults to the file-system-type named “sanfs”. If there is no such type in the /etc/vfs file, the system defaults to the file-system-type number 20.

You will use this name to create the virtual client.

–fstypenum *number*
Identifies the number associated with the file-system type for the file-system-driver instance. All mount requests for a file system of this type are routed to this file-system-driver instance.

You would use this parameter only when you load multiple instances of the file-system driver on the same client system.

–debuginit
Enables the file-system driver to issue diagnostic messages of the CONFIG class. Messages in this class are issued only during initialization.

If specified, the file-system driver does not issue diagnostic messages. You can turn the messages on after the file-system driver is running using the **stfsdebug** command.

Note: This parameter is intended for use only by trained service technicians.

–tracefile *file_name*
Specifies that the file-system driver is to write diagnostic information to the specified file.

Note: The specified file must already exist.

–synctrace
Specifies that the file-system drive is to write diagnostic information to the specified trace file synchronously rather than using buffered writes.

driver_file_name
Specifies the name and location of the file-system driver that you want to load, unload, or query. The file name is typically "sanfs".

The file-system driver is loaded as a kernel extension. To identify the instance of the file-system drive, you identify the kernel extension. The kernel-extension name is the same as name and location of the file-system driver that was used to load the driver (for example, /usr/tank/client/bin/sanfs).

Description:

Prerequisite: You must have root privileges to use this command.

This command creates a file-system-driver instance by loading the file-system driver as a kernel extension. This command also unloads or queries the kernel extension.

After loading the file-system driver, you can use the **stfsclient** command to create a virtual client and then use the **stfsmount** command to mount the global namespace.

Examples: Loads the file-system driver The following example loads the file-system driver on a client for AIX.

```
stfsdriver –load /usr/tank/client/bin/sanfs
```

Related topics:

- "Clients" on page 10
- "stfsclient" on page 303
- "stfsmount" on page 311
- "UNIX-based clients" on page 12

stfsmount

Mounts the global namespace.

Syntax:

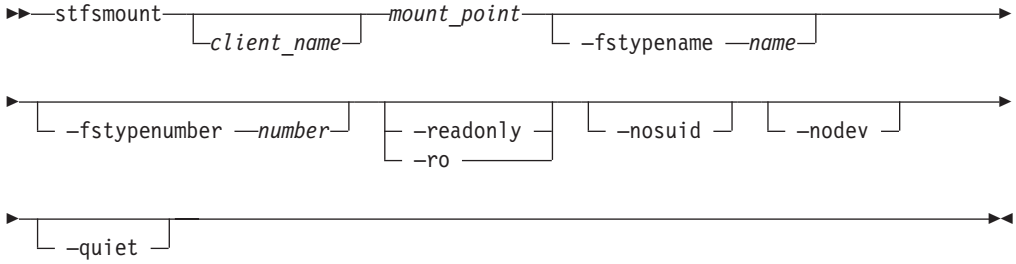

Parameters:

client_name
Identifies the unique name of the virtual client to which you want to mount the global namespace. The client must be up and running. The default client name is the host name of the client system.

mount_point
Specifies the directory associated with the global namespace image that you want to mount.

–fstypename *name*
Specifies the name of the file-system type to use for the file-system-driver instance. This is the same name used to load the file-system driver.

This name relates to a specific file-system-type number. The file /etc/vfs maps the file-system-type name to the number.

If you do not specify a file-system-type name or number, the system defaults to the file-system-type named “sanfs”. If there is no such type in the /etc/vfs file, the system defaults to the file-system-type number 20.

You would use this parameter only when you load multiple instances of the file-system driver on the same client system.

–fstypenumber *number*
The number that identifies the file-system type for the file-system-driver instance. All mount requests for a file system of this type are routed to this file-system-driver instance.

–readonly | –ro
Sets the global namespace image to read only. If specified, an attempt to update data or metadata in the global namespace will fail, and an attempt to access a file-system object will not update its access-time attribute.

–nosuid
Disallows any invocation of the setuid or setgid commands from this file-system image.

–nodev
Disallows any attempts to open device nodes in this file-system image.

–quiet
Specifies not to display informational messages. This parameter does not affect error messages.

Description:

Prerequisite: You must have root privileges to use this command.

This command creates an image of the global namespace on the client system by mounting a directory. The global namespace maintains a list of its directories that are available to the clients. When a client mounts a directory in the global namespace, that directory and its subdirectories become part of the client's directory hierarchy.

Note: This command is used in place of the **mount** command to mount the global namespace.

Before you can mount the global namespace, you must have a virtual client running on the client system. To create the virtual client, use the **stfsclient –create** command.

Remounting the global namespace image is not the same as unmounting the global namespace and then mounting it again. Rather, it changes the attributes of an existing global namespace image, such as changing from read-write to read-only mode. To remount the global namespace image or to see what global namespace images currently exist, use the **stfsmount** command.

To unmount the global namespace, use the **stfsumount** command.

Examples: Mount the global namespace The following example mounts the global namespace.

```
stfsmount mnt/SANFS_MOUNTPT –fstypename sanfs
```

Related topics:

- "Clients" on page 10
- "UNIX-based clients" on page 12
- "stfsclient" on page 303
- "stfsumount" on page 314

stfsstatus

Displays the version of the file-system driver for the specified virtual client for AIX.

Syntax:

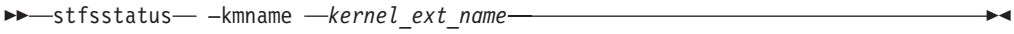

Parameters:

–kmname *kernel_ext_name*
Identifies kernel-extension name of the file-system driver associated with the virtual client.

The file-system driver is loaded as a kernel extension. To identify the instance of the file-system drive, you identify the kernel extension. Each kernel extension has a name, but this name is not unique. This name is usually the file name of the object file from which you loaded the kernel extension (for example, /usr/tank/client/bin/stfs). To determine the kernel-extension name, use the **genkex | grep stfs** command.

Description:

Prerequisite: You must have root privileges to use this command.

After issuing this command, if the client is running, the version of the file-system driver is displayed. If the file-system driver is not loaded, an error message is displayed stating that system could not determined the file-system driver instance.

Examples: Display the file-system-driver version The following example displays the version of the file-system driver for the local client.

```
stfsstatus –kmname /usr/tank/client/bin/stfs
```

Related topics:

- "Clients" on page 10
- "stfsclient" on page 303
- "stfsdriver" on page 309
- "UNIX-based clients" on page 12

stfsumount

Unmounts the global namespace.

Syntax:

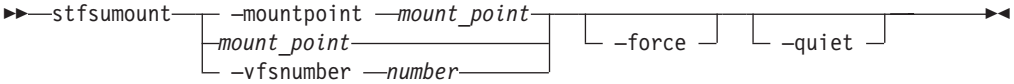

Parameters:

–mountpoint *mount_point* | *mount_point*
Specifies the directory associated with the global namespace image that you want to destroy. This must be the same directory that you specified to mount the global namespace.

If you created multiple global namespace images over the same directory (mount point), this command chooses the most recently created directory.

–vfsnumber *number*
Identifies the virtual file-system (VFS) number associated with the global namespace image that you want to destroy.

In AIX, every global namespace image has a unique VFS number. The **stfsmount** command displays this number when it creates the global namespace image.

Note to reviewers:

What VFS stand for virtual file system?

–force
Unmounts the file system even if it is in use.

–quiet
Specifies not to display informational messages. This parameter does not affect error messages.

Description:

Prerequisite: You must have root privileges to use this command.

This command destroys a global namespace image on the client system. It is used in place of the AIX **umount** command.

After you destroy all global namespace images that are linked to a specific virtual client, you can destroy the virtual client using the **stfsclient –destroy** command.

To see what global namespace images currently exist, use the AIX **mount** command with no parameters.

Examples: Unmount the global namespace The following example unmounts the global namespace on the local client.

```
stfsumount –mountpoint sanfs/cnt1
```

Related topics:

- "Clients" on page 10
- "UNIX-based clients" on page 12
- "stfsmount" on page 311

Windows-client command

The following table provides a brief description for the Windows-client command.

Note: You must have Administrator privileges on the Windows client to use these commands.

Command	Description
Migration	
migratedata	Migrates data to SAN File System.

migratedata

Migrates data to SAN File System.

Syntax:

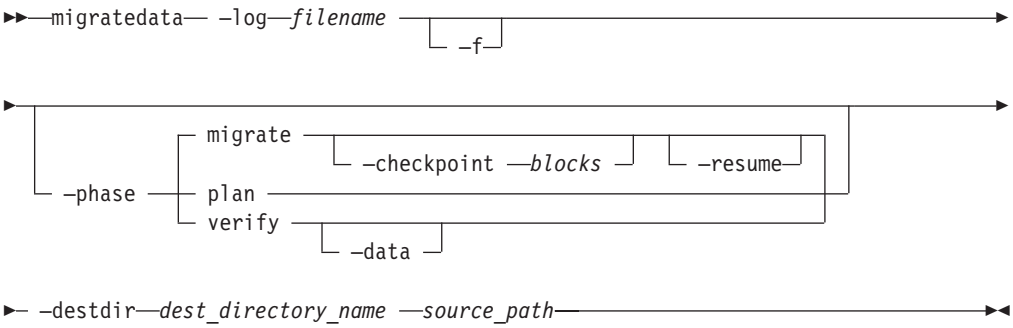

Parameters:

–log *filename*
Specifies the location of a file in which to log migration activities, warnings, and errors. When used with the **–plan migrate –resume** parameter, the **–log** parameter specifies the log file from which to read information about the last completed block or file.

Attention: You must specify the correct log file with migrate –resume and verify that the source and destination directories specified on the command line match those in the log file.

- If you specify an incorrect log file and the **–f** parameter, –resume displays a warning and overwrites the target file-system data with wrong information.
- If you specify an incorrect log file, but do not specify the **–f** parameter, this command displays an error and exits.

–f Specifies that the migration should continue even if there is an error with a file. If specified with the **–phase migrate** parameter, this command skips any files with errors, but continues with the migration process. If not specified, an error results in the entire migration being stopped before the file that caused the error. You can then restart the migration after fixing the error.

If specified with the **–phase verify** parameter, this command adjusts any missing metadata attributes, such as permissions and times. If there is a mismatch in size, however, this command will not remigrate the file.

–phase
Specifies the migration phase to run. Choices include:

plan Gathers information about the available system resources (available memory, number of CPUs, size of the source tree and space available on the destination file system), copies sample files from source directory to estimate transfer rates, and provides an estimated time for the migration of the data set.

migrate
Reads data from source file system and writes the data to the destination file system. Although not required, for large data sets, you should run this command in planning mode first. You can stop the migration process at any point and resume from the last completed file or block (using the **–resume** parameter).

This is the default value.

verify Verifies the integrity of the migrated data using the Message Digest 5 verification algorithm on the contents of the file, as well as verifying consistency of the metadata (such as owner and modification time stamp settings) between the source and destination files.

You can specify more than one phase. For example, to plan, migrate, and verify the data, specify **–phase plan –phase migrate –phase verify**. Although you can specify the phases in any order, this command always estimates the completion time, migrates data, and then verifies the migrated data.

If the **–phase** parameter is not specified, this command runs only the migration phase.

–checkpoint *blocks*
Writes a checkpoint in the log file after each specified number of blocks of file data has completed the migrate phase. (The block size depends on the client platform.) For example, if you specify **–checkpoint 20**, this command makes an entry in the log file each time 20 blocks of file data is migrated. On a platform with a block size of 16 MB, this command writes to the log file after each 2 560-MB of the file data has been migrated. If the process is interrupted, you can resume the migration at the place it left off. If unspecified, the **migratedata** command makes an entry in the log file after each complete file has been migrated. You can resume the migration at the point of the last migrated file.

–resume
Resumes the migration from the last completed block or file (logged in the log file specified by –log). If the log file indicates that some files in the source directory are migrated and this parameter is not specified, this command restarts the migration process from the beginning (performs a fresh migration).

–data
Verifies every block of source data (file data and metadata) with the destination data. If not specified, this command verifies only the metadata unless there is a mismatch in the file attributes, in which case this command then verifies the file data.

Note: Verifying all data is very time consuming and can take as long as the migration itself.

–destdir *dest_directory_name*
Specifies the name of the destination directory for the migrated data. The directory can either exist or be a new directory name. IBM recommends that you create the directory before beginning the migration process. If the directory does not exist, this command creates the directory using the default permissions.

source_path
Specifies one or more paths of directories or files to migrate.

Description:

Prerequisite:

1. You must have root privileges on AIX or Administrative privileges on Windows to use this command.
2. All storage pools, all filesets, and at least one policy must be set up. All activity (from applications, such as database servers and application servers, or users)

that modifies data on the source and destination file systems must be stopped and remain stopped to guarantee consistency of the migrated data.

3. The destination directory must exist with correct set of permissions and appropriate storage policies must be configured.

Examples: Migrating data from a client for AIX This example migrates data from the work/capital directory on the client machine to the sanfs/cnt1 directory in the global namespace. A checkpoint is written to the mgrt_capital.log log file each time 20 blocks of file data is migrated.

```
migratedata –log /mgrtlogs/mgrt_capital.log –phase migrate –checkpoint 20
–destdir /mnt/tank/sanfs/cnt1 work/capital
```

Migrating data from a client for Windows This example migrates data from the C:\Capital directory on the client machine to the sanfs\cnt1 directory in the global namespace. A checkpoint is written to the mgrt_capital.log log file each time 20 blocks of file data is migrated.

```
migratedata –log c:\mgrtlogs\mgrt_capital.log –phase migrate –checkpoint 20
–destdir t:\cnt1 C:\Capital
```

Service commands and utilities

Administrative commands:

The following table provides a brief description and role for each command in the Administrative CLI that is intended for use only by trained service personnel.

Command	Description	Role	Environment
Administrative server			
startCimom	Starts the Administrative agent.	root	shell
stopCimom	Stops the Administrative agent.	root	shell
Cluster			
chadvclusterconfig	Modifies the cluster settings that require a reboot when changed.	Administrator	tanktool
chclusterconfig	Modifies the cluster settings that do not require a reboot when changed. **Note:** This command may be used by administrative users. Only specific options in this command are intended for use only by trained service technicians.	Administrator	tanktool
mktruststore	Creates a truststore that is shared by the administrative infrastructure to certify secure connections.	n/a	shell
tank extractbootrecord	Extracts a local Tank.Bootstrap file from a valid master volume.	Administrator	shell
tank lscluster	Displays the cluster definition from a system master volume when the Metadata servers are not running.	Administrator	shell
tank lsdisklabel	Displays the a SAN File System product label for a specified device.	Administrator	shell
tank lsversion	Displays the version control information from a system master disk.	Administrator	shell
tank resetcluster	Erases the static cluster definition contained on the system master volume.	Administrator	shell

Command	Description	Role	Environment
tank resetversion	Resets the version-control information on a system master disk.	Administrator	shell
Engine			
pmf	Collects information for the first-failure data capture of SAN File System problems from the engine on which the Metadata server resides.	Operator, Administrator	shell
Legacy CLI			
legacy	Runs commands in the legacy CLI.	Administrator	shell
Metadata server			
addserver	Adds a Metadata server on a specific engine to the cluster.	Administrator	tanktool
dropserver	Stops and drops a Metadata server on a specific engine from the cluster.	Administrator	tanktool
SAN File System console			
disableConsoleTrace	Disables tracing and logging for the SAN File System console.	root	shell
enableConsoleTrace	Enables tracing and logging for the SAN File System console.	root	shell
startConsole	Starts the SAN File System console Web server.	root	shell
stopConsole	Stops the SAN File System console Web server.	root	shell

Client commands:

The following table provides a brief description and role for each client command that is intended for use only by trained service personnel.

Command	Description	Role
Diagnostics		
pmf	Collects information for the first-failure data capture of SAN File System problems from the client machine.	root (AIX) or Administrator (Windows)
stfsdebug	(AIX only) Enables or disables the logging of file-system-driver debug messages in the syslog for the specified virtual AIX client.	root

Command	Description	Role
stfsdriver	(AIX only) Loads the file-system driver as a kernel extension. **Note:** This command may be used by administrative users. Only specific options in this command are intended for use only by trained service technicians.	root
stfsstat	(AIX only) Display statistics about the client.	root

addserver

Adds a Metadata server on a specific engine to the cluster.

Syntax

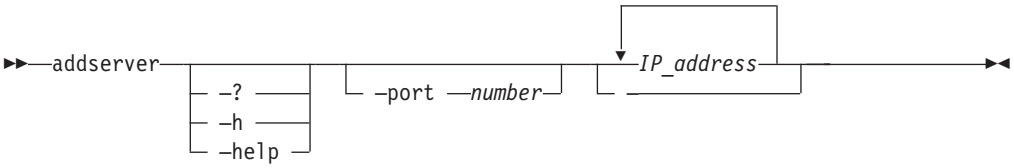

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–port *number*
Specifies the port number on the engine to which to add the Metadata servers. The default value is the default cluster port, 1737.

IP_address
Specifies one or more Internet Protocol (IP) addresses associated with the Metadata server to add.

– Specifies that you want to read one or more IP addresses that are associated with the Metadata server to add from stdin (for example, – << /work/server_list.txt).

Description

Attention: This task must be performed only by trained service technicians.

Prerequisites:

1. This task must be performed only by trained service technicians.
2. You must be logged in to the engine that hosts the master Metadata server to run this command.

Typically, you would issue this command during set up. You must run this command from the engine that is running the master Metadata server. After you add the Metadata servers to the cluster, they can begin serving the workload.

Examples

Adds a Metadata servers to the cluster The following example adds the Metadata server on the engine 179.24.17.34 to the cluster.

```
tanktool> addserver 179.24.17.34
Server 179.24.17.34 successfully added to the cluster.
```

Related topics

- "Metadata server" on page 8
- "dropserver" on page 326

chadvclusterconfig

Modifies the advanced cluster settings.

Syntax

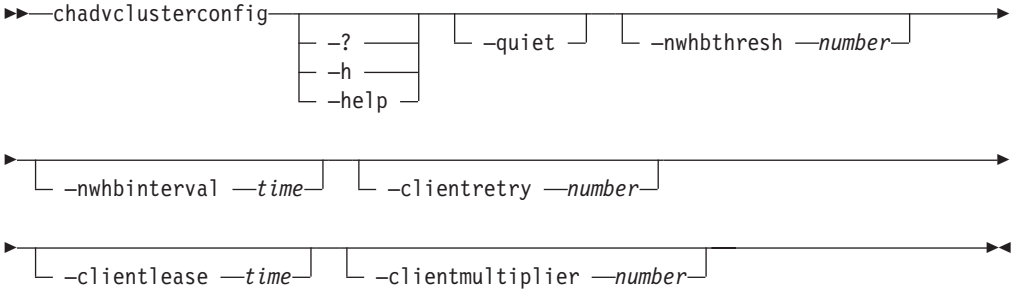

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–quiet
Turns off the prompt to confirm the change operation

–nwhbthresh *number*
Specifies the maximum number of missed network heartbeats before the Metadata server declares a possible path-failure problem with another Metadata server in the cluster. If this threshold is reached or exceeded, the Metadata server goes into the not-running state, if it was up, and stays in that state until it is repaired. Otherwise, the ejected Metadata server would crash, restart, and rejoin the cluster.

You can specify a value from 1 to 100. The default value is 40.

–nwhbinterval *time*
Specifies the network heartbeat interval, in milliseconds, for periodic pings among Metadata servers in the cluster. You can specify a value from 200 to 10 000. The default value is 500.

–clientretry *number*
Specifies the maximum number of times a send can be attempted from the Metadata server to the client before declaring the client's session lease to be expired. You can specify a value from 1 to 100. The default value is 5.

–clientlease *time*
Specifies the amount of time, in seconds, that a lock is leased to a client when the Metadata server grants a lock. The Metadata server applies the client-lease-period multiplier to this value before expiring the lease. You can specify a value from 10 to 120. The default value is 20.

–clientmultiplier *number*
Specifies the client-lease-period multiplier. This multiplier accommodates possible networking delays between the time a message gets sent from the client until it is received by the Metadata server. You can specify a value from 0 to 4. A value of 0 prevents this multiplier from affecting the client-lease period. The default value is 2.

Description

Attention: This task must be performed only by trained service technicians.

Prerequisites: This task must be performed only by trained service technicians.

Note: Modifying these settings requires you to reboot the cluster before changes will apply.

After you run this command, you cannot run the **chadvclusterconfig** or **chclusteconfig** command until you restart the cluster.

After you modify any of these disruptive cluster settings, you cannot modify any cluster settings (using the **chadvclusterconfig**or **chclusteconfig** command) until you restart the cluster.

Examples

Changes the network heartbeat threshold The following example changes the network heartbeat threshold to 10 missed heartbeats.

```
tanktool> chadvclusterconfig -nwhbthresh 10
Are you sure you want to change advanced cluster configuration settings,
which will require a immediate cluster restart for changes to apply? [y/n]  y
Cluster successfully modified. Restart the cluster for all changes to apply.
```

Related topics

- "chclusterconfig" on page 172
- "startcluster" on page 272
- "statcluster" on page 278
- "stopcluster" on page 290
- "upgradecluster" on page 295

disableConsoleTrace

Disables tracing and logging for the SAN File System console.

Syntax

Description

Prerequisites::

1. This task must be performed only by trained service technicians.
2. You must have root privileges to use the command.
3. The Web server must be running.

Note: This command is run from the shell prompt. It is not run inside of tanktool.

This command is case sensitive and must be entered as shown.

Tracing is automatically disabled after you restart the Web server.

Examples

Disable tracing The following example disables tracing and logging for the SAN File System console.

```
#disableConsoleTrace
Disabling SAN File System console tracing...
```

Related topics

- "enableConsoleTrace" on page 327
- "startConsole" on page 344
- "stopConsole" on page 352

dropserver

Stops and drops a Metadata server on a specific engine from the cluster.

Syntax

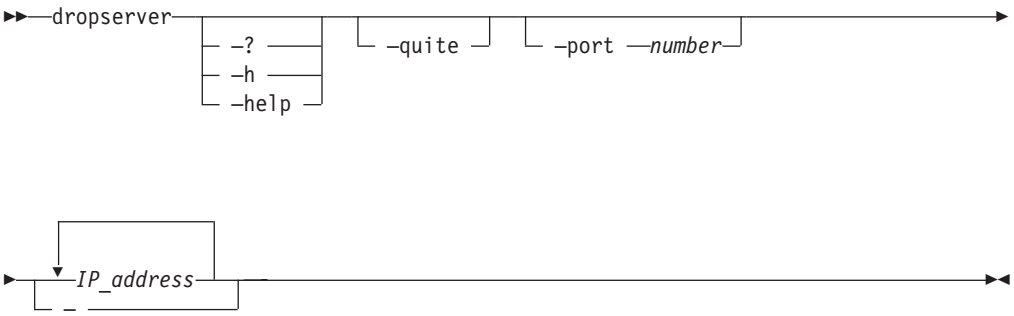

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

–quiet
Turns off the prompt to confirm the remove operation.

–port *number*
Specifies the port number on the engine from which to drop the Metadata servers. The default value is the default cluster port, 1737.

IP_address
Specifies one or more IP addresses associated with the Metadata server to drop.

– Specifies that you want to read one or more IP addresses associated with the Metadata server to drop from stdin (for example, – << /work/ipaddr_list.txt).

Description

Prerequisites:

1. This task must be performed only by trained service technicians.
2. You must be logged in to the engine hosting the master Metadata server to run this command.

This command is typically issued when changing the IP address or upgrading the Metadata server. You cannot drop the last Metadata server in the cluster.

Examples

Drops a Metadata servers from the cluster The following example drops the Metadata server on the engine 179.24.17.34 from the cluster.

```
tanktool> dropserver 179.24.17.34
Server 179.24.17.34 dropped from the cluster.
```

Related topics

- "Metadata server" on page 8
- "addserver" on page 322

enableConsoleTrace

Enables tracing and logging for the SAN File System console.

Syntax

Description

Prerequisites::

1. This task must be performed only by trained service technicians.
2. You must have root privileges to use the command.
3. The Web server must be running.

Note: This command is run from the shell prompt. It is not run inside of tanktool.

This command is case sensitive and must be entered as shown.

Tracing is automatically disabled after you restart the Web server.

This command save log and trace records to the /opt/was/logs/server*x*/trace.log file.

Enabling tracing will impact the performance of SAN File System. Use this command only when necessary.

Examples

Enable tracing The following example enables tracing and logging for the SAN File System console.

```
enableConsoleTrace
Enabling SAN File System console tracing...
```

Related topics

- "disableConsoleTrace" on page 325
- "startConsole" on page 344
- "stopConsole" on page 352

legacy

Runs commands in the legacy command-line interface.

Syntax

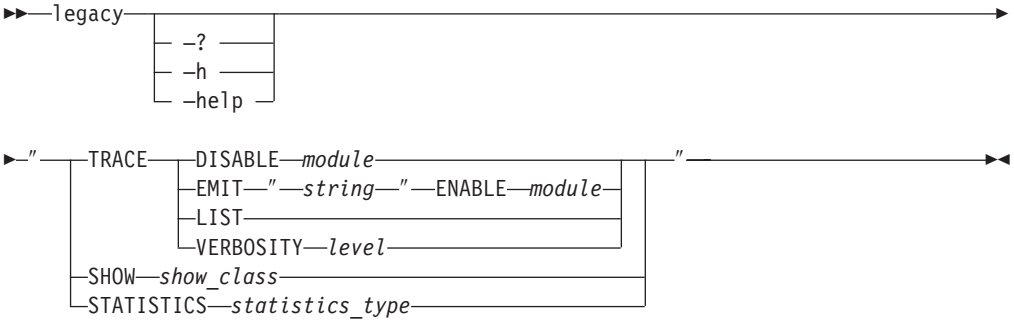

Parameters

–? | –h | –help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

TRACE
Configures tracing functions. If this parameters is specified without any additional options, this command prints online help about all the trace command variants

DISABLE *module*
Disables tracing on the specified module. If you do not specify a module, this command prints out all of the disabled trace modules. You can use single-wildcard (?) and multiple-wildcard characters (*) to specify more than one module (for example, TM:* would disable all of the TM modules, and WAL:WALF????:* would disable theWAL:WALFLUSH:WALWRITERREP and WAL:WALFORCE:WALWRITERREP modules.

EMIT *"string"*
Write the specified string to the trace log, , /usr/tank/server/log/log.trace.

ENABLE *module*
Enables tracing on the specified module. If you do not specify a module, this command prints out all of the enabled trace modules. You can use single-wildcard (?) and multiple-wildcard characters (*) to specify more than one module (for example, TM:* would enable all of the TM modules, and WAL:WALF????:* would enable theWAL:WALFLUSH:WALWRITERREP and WAL:WALFORCE:WALWRITERREP modules.

LIST Displays a list of all the trace modules that you can enable or disable in the Metadata server. If a module is enabled, it is prefixed with an asterisk (*).

VERBOSITY *level*
Sets amount of output that is generated from tracing, based on the verbosity level. You can specify a value from 0 to 9. A value of 0,

which is the default, emits only the most important messages, whereas a value of 9 emits all messages. If you do not specify a level, the current verbosity level is displayed.

SHOW *show_class*
Displays information about the specified show class. You can specify one of the following show classes:

ASMSESSIONS
Displays information about the administrative sessions.

BUFPOOL [*master*]
Displays information about pages in the buffer pool, both the clean list and dirty list. The *master* parameter is a flag that can be set to 0 or 1. The default is 0.

If the master Metadata server is not specified, this command displays the buffer pool of the subordinate Metadata server's workload. Specifying SHOW BUFPOOL 1 displays the buffer pool of the master Metadata server's workload.

CLEANLIST [*master*]
Displays information about only clean pages in the buffer pool (see BUFFPOOL).

CMATTACHPOINTS | CMATT
Displays the attach-points table.

CMARENAS
Displays the arenas table.

Note: The output for this command can be quite large.

CMCONTAINERS | CMCONT
Displays information about the filesets (or containers) served by the subordinate Metadata server.

CMLOADMAP [*detail*] **| CMLOAD** [*detail*]
Displays the load map, which maps the write-ahead log to the Metadata server. If detail > 0, displays additional information.

CMSMASTERREC | CMM
Displays the contents of the Cluster Manager master record.

CMSECTOR *sector_number*
Reads the raw contents of a Cluster Manager sector *sector_number* from LVM and displays the data in "dump" form.

CMSPACERECLAIM | CMSPREC
Displays information about the space-reclamation thread on the master Metadata server.

CMSTGPOOLS | CMSTG
Displays information about the storage pools.

CMSVCCLASSES | CMSVC
Displays the service-class table.

CMVOLUMES | CMVOL
Displays the volumes table.

CMCONTBIND | CMCB
Displays the fileset-bindings table, which maps the write-ahead-log volume each fileset is bound to.

CMWALVOLS | MWALS
Displays the write-ahead-log volumes table.

CONDITION | COND
Displays information about the condition of the Metadata server.

DBPAGE *space_ID page_address* **[***format***] [***master***]**
Displays the contents of a page.

DBPAGEHDR *space_ID page_address* **[***format***] [***master***]**
Displays the contents of a page (header only).

DBPAGETABLE [*master***]**
Displays information about locked pages (see BUFPOOL).

DBSPACE [*master***]**
Displays information about DB Spaces (see BUFPOOL).

DBTXNTABLE [*master***]**
Displays information about in-flight transactions (see BUFPOOL).

DIODISKS
Displays the open disk table.

DIRTYLIST [*master***]**
Displays information about only dirty pages in the buffer pool. (See BUFPOOL)

DISPATCHERSTATS | DSPSTATS
Displays Dispatcher and queue statistics.

FLASHCLEANER *fileset_ID* **[***flags***] [***dump_ID***]**
Displays information about a fileset's (container's) FlashCopy Images.

FLASHTABLE [*fileset_ID***]**
Displays information about a fileset's (container's) FlashCopy Image table.

FSCK Displays information about an active metadata check that is in progress.

FSCKSUB
Displays information about an active metadata check that is in progress on a subordinate Metadata server.

GIODISKS
Displays information about the global disk.

GSMSGSESSION

GSNODES
Displays information about the engine.

HAMSTATE

INDEXSTATS *space_id page_address* **[***format***] [***master***] | INDEX** *space_id page_address* **[***format***] [***master***]**
Verifies the structural integrity of an index, and shows statistics on the index as a side-effect.

LATCH
Displays all active latches.

LATCHSTATS
Displays latch contention statistics.

LATCHX
Displays all active exclusive latches.

LATCHXSTATS
Displays exclusive-latch contention statistics.

LMCLIENTS
Displays a list of clients that are known to the lock manager.

LMALLLOCKS
Display all of the locks in the Lock Manager.

LMLOCKSBYCLIENT [*client_ID***]**
Display all of the locks held by the specified client. To display a list all registered clients, use TMCLIENTS.

LMLOCKSBYOID [*O***] [***I***] [***D***] | LMLOCKS [***O***] [***I***] [***D***]**
Display all of the locks held for the specified object.

LMSTATS
Displays summary statistics on the lock manager.

LVMDISKS
Displays the disk table.

LVMMASTERREC | LVMMR
Displays master control block, including:
- Global ID of master disk.
- Sector size of master disk.
- Size, in bytes, of the logical and physical partitions for this Logical Volume Manager installation.
- Installation common sector size, bytes.
- Number of sectors per partition.
- Number of physical partitions reserved to hold Logical Volume Manager persistent metadata.
- Starting sector numbers for the shadow copies of the Logical Volume Manager persistent tables.
- Index (0 or 1) of the shadow copy that contains the committed copy of the Logical Volume Manager persistent tables
- Update sequence (version) number.

LVMPARTMAP | LVMPM *volume_ID*
Displays a volume's partition map.

LVMREP
Displays information about the volume manager on a subordinate Metadata server.

LVMVOLDESC *volume_ID*
Displays information about a volume's description.

LVMVOLUMES
Displays the volume table, including:
- Logical volume class.
- Logical volume ID.
- Size, in bytes, of logical pages.
- Capacity, in number of logical partitions.
- Number of formatted LPs.
- Node ID of subordinate that has locked this logical volume.
- Partition map for each logical partition, including
  - Logical Partition.
  - LVM-assigned disk number of physical disk.

- Partition status flags.
- Physical partition number on disk.

MASTERWALSTATS | MWALSTATS
Displays write-ahead-log statistics for the master Metadata server.

MUTEX
Displays all active mutexes.

MUTEXSTATS
Displays statistics about mutexes.

OMCONTAINERS | OMCONT [*fileset_ID***]**
Displays all fileset (or container) information.

OMCONTAINERSTATS | OMCS [*fileset_ID***]**
Displays statistics about a specific fileset (or container).

OMOBJ [raw | pit | write | main] [*object***] [count** *count***] [start** *start***]**
Displays object.

OPTCONFIG
Displays information on the Metadata server configurable parameters.

PMTABLE
Displays information about long-running administrative commands.

SCARENASTATS | SCAS [*fileset_ID***] [***pool_ID***]**
Displays statistics about a fileset's (or container's) free-space map.

SCSTSDSTATS | SCSS [*fileset_ID***]**
Displays statistics about a fileset's (or container's) STSD table.

SERVERHANG
Displays the major data structures used for detecting when a Metadata server hangs.

STATEINFO
Displays information about the state of the cluster and Metadata servers.

SUBWALSTATS
Displays write-ahead-log statistics for the subordinate Metadata server.

THREAD (STATE [force]) | (STACK [ALL | *thread_ID*]) | USAGE | HELP
Displays information about threads.

TMATIMESTATS
Displays information about tmatime statistics.

TMATIMESTATSRESET
Resets information about tmatime statistics.

TMCLIENTS
Displays statistics for registered clients.

TMFILTERS
Displays information about Metadata server transaction-manager filters.

TMPATH *cluster_ID fileset_ID local_OID* **[***epoch_ID***]**
Displays information about Metadata server transaction-manager filters.

VERSION
Displays information about the Metadata server version.

WALCKPT [*master***]**
Displays write-ahead-log context for the local node (see BUFPOOL).

WALSTATS [*master***]**
Displays write-ahead-log statistics (see BUFPOOL).

WALWRITER [*master*]
Displays write-ahead-log statistics (see BUFPOOL).

STATISTICS *statistics_type*
Displays information about the specified show class. You can specify one of the following types:

RESET
Resets the statistics buffer.

SHOWALL
Displays all statistics about the CM master transactions and the SAN File System protocol received by the master and subordinate Metadata servers.

SHOWBTREE
Displays BTree statistics.

SHOWBUF
Displays statistics about the buffer pool.

SHOWCM
Displays statistics about the CM master transactions received by the master and subordinate Metadata server.

SHOWDSP
Displays dispatcher and queue statistics.

SHOWGIO
Displays statistics about the global disk table.

SHOWLATCH
Displays statistics about latches.

SHOWLATCHS
Displays statistics about extended latches.

SHOWLM
Displays statistics about the Lock Manager.

SHOWMUTEX
Displays statistics about mutexes.

SHOWSTP
Displays statistics about the SAN File System protocol by message type received by the Metadata servers.

SHOWWAL
Displays write-ahead-log statistics for the master and subordinate Metadata server.

Description

Prerequisite:: This task must be performed only by trained service technicians.

If the legacy command string contains single quotation marks (′) or double quotation marks (″), you must proceed the character with a backslash.

Examples

Running a legacy command The following example runs the show command from the legacy CLI.

```
tanktool> legacy "SHOW TMCLIENTS"
List of identified clients:
  ClientId=1010513811, ClientName="st-client-1", SendInProg=0, SendMsgNo=2,
  ccb_sendAcked=1, DeliveredMsgNo=1, IPaddress=192.168.1.10, IPport= 32769
```

mktruststore

Obtains a public certificate from LDAP to and creates a truststore that is shared by the administrative infrastructure to certify secure connections. This command replaces any existing truststore on the local engine.

Syntax

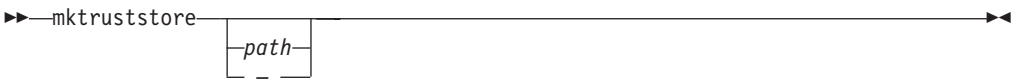

Parameters

path

Specifies the full directory path to import an existing LDAP certificate into the new truststore file being created. The certificate cannot be added to an existing truststore because a new truststore will be created. If not specified, an LDAP certificate will not be added to the new truststore.

Note: An LDAP certificate is required to obtain secure communication with the LDAP server.

– Specifies that you want to read the path from stdin (for example, – << /work/path.txt).

Description

Prerequisites:

1. The cimom.properties file must be set up prior to using this command.
2. The tank.properties file must exist. It is used by the mktruststore command to determine the language being used.
3. This task must be performed only by trained service technicians.

Note: This command is run from the shell prompt from the /usr/tank/admin/bin directory. It is not run inside of tanktool.

The truststore file resides in the /usr/tank/admin/directory.

You would use the **mktruststore** command is used in these circumstances:

- During installation, the **mktruststore** command is run automatically by the **setupTank** script on the first engine. Because the truststore file must be exactly the same on every engine in the cluster, you must copy the truststore file from one engine to each remaining engine in the cluster before running setupTank on those engines.
- When replacing an expired truststore file. The truststore file is valid for one year.
- When you need to change the truststore (for example, if security is breached). You may change the truststore at any time; however, it must be the same on every engine in the cluster. The Administrative agent must be restarted any time the truststore is changed by issuing the **stopcimom** and **startcimom** commands. The user interfaces (SAN File System console and Administrative command-line interface) connection will be broken when the Administrative agent is stopped.

Examples

Create truststore The following example creates a truststore and imports the LDAP certificate ldap.cert from the local directory.

```
mktruststore ldap.cert
Creating truststore file.
The truststore was created successfully
Certificate was added to keystore.
```

Related topics

- "Administrative server" on page 8
- "startCimom" on page 343
- "stopCimom" on page 351

pmf

Collects information for the first-failure data capture of SAN File System problems, from either a client machine or the storage engine.

Syntax

Linux systems and clients for AIX

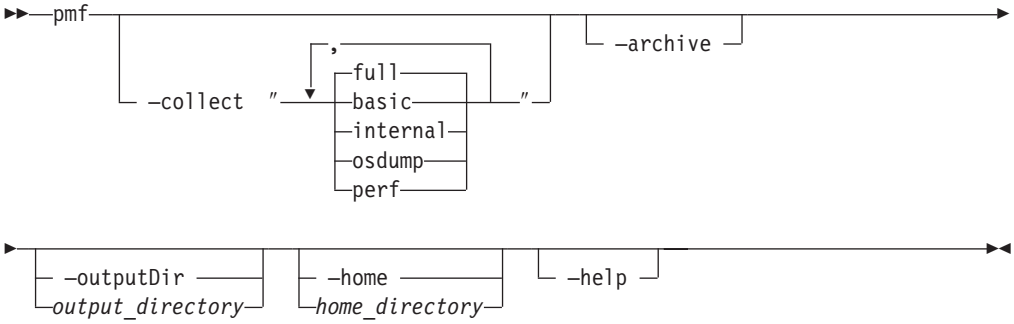

Clients for Windows

Parameters

–collect basic,full,internal,osdump,perf

Specifies the type and amount of data to be collected. You can specify more than one value, separated by a comma without spaces between. Each specified value must be enclosed in double quotes on Windows or in double or single quotes on UNIX.

basic Collects only a small amount of fundamental system data, which is useful for addressing simple usage and configuration questions.

full Collects all of the data necessary to handle the majority of problems. If not specified, this command gathers the full amount of data by default.

Note: The **full** and **basic** values are mutually exclusive. If both are specified, the **full** value takes precedence.

internal

Collects thread stacks for a Metadata server process, which are useful when more detailed diagnostic data is needed to solve a difficult problem. The internal data is not gathered with the **full** value. (The **internal** value is currently undefined and ignored for a client.)

osdump

Collects a kernel memory image and the operating-system binary file that generated the image. This information is most useful for investigating problems resulting in the crash of the operating system on a client machine, possibly as a result of a problem with the installed SAN File System kernel module. This information is of less value for a Metadata server because the Metadata server process runs in user-space. However, this value will cause the **pmf** command to collect an operating system binary file and a core file it has generated from a Metadata server machine.

Operating system core files are not gathered with the **full** value, though use of that value does cause the command to collect a core file generated by a Metadata server or Administrative server process.

Note: Note that use of the **osdump** value can cause the **pmf** command to collect a very large amount of data. For example, a client machine with 2 GB of memory will produce a core file of the same size. (The binary file for the operating system can also be somewhat large). Therefore, you must take steps to ensure availability of sufficient disk space in the output directory before invoking the **pmf** command with this value.

perf Gathers statistics and metrics about the performance of a Metadata server or client, which is useful in cases where system response has degraded. The performance data is not gathered with the **full** value.

–archive
Compresses and packages all the collected data into a single archive or zip file, named SANFS_pmf_archive, that can be sent from the customer site to IBM support personnel, usually using the **ftp** utility. On UNIX platforms, the **tar** command is used to assemble the single package; the **compress** command is used to compress the data on AIX, and the **gzip** utility is used (through the **–z** parameter of the **tar** command) for data compression on Linux. On the Windows 2000 platform, the command uses the **pkzip** or **winzip** command (whichever is listed as the default application in the Windows Registry) to assemble and compress the data. By default, the **pmf** command performs no archiving or compression of the data it collects, though you can perform these operations on the collected data after the command has executed.

–outputDir *output_directory*
Specifies the directory in which to assemble the collected. By default, collected data is gathered in subdirectories of the directory indicated by $TANKDIR/pmf on the local machine on which it physically executes. You can use this parameter to specify an alternative output directory as a full or relative path name; a relative pathname is concatenated to the directory from which this command is invoked.

–home *home_directory*
Specifies the root of the installation directory for the SAN File System component installed on the local machine. By default, the command relies on the value of the $TANKDIR environment variable to determine the location at which SAN File System is installed. The user can specify the home directory as a full or relative path name; a relative path name is concatenated to the directory from which this command is invoked.

–help
Displays the syntax of this command. If specified, all other parameters specified with the command are ignored.

Description

Prerequisite: This task must be performed only by trained service technicians.

The problem-monitoring-facility (pmf) command gathers data for diagnosing errors or failures associated with Metadata servers and clients. This command is intended primarily for first-failure data-capture capabilities useful for investigating problems upon their initial occurrence, without requiring problem recreation or subsequent tracing. The **pmf** command is also referred to as the one-button data collection because it collects problem-determination data with a single invocation.

This command must be invoked natively on the Metadata server or client machine from which you want to collect information. For a Metadata server, this command can be invoked using the SAN File System console or Administrative command-line interface (CLI). The CLI command is named **collectdiag**. For a Metadata server or client, the command can be invoked from a command shell (on UNIX) or a command prompt (on Windows). The SAN File System console provides true one-button execution of the command against a specified Metadata server, but it always issues the command with the default parameters for all parameters. The command must be invoked using one of the alternate methods if non-default parameters are needed.

Linux systems and clients for AIX

On the storage engine and client machines running AIX, each invocation of the **pmf** command results in creation of two subdirectories nested under the output directory. The name of the first subdirectory is based on the short form of the host name of the machine on which the command runs. This directory allows output from multiple machines to be identified clearly when output is written to the same remote device. The name of the second subdirectory, which is created within the first, is based on the date and time at which the command was run. Timestamp directory preserves data from multiple invocations of this command for the same machine. For example, if you invoke the pmf command from a machine named sanfs.ibm.com at approximately 8:05 AM on 9 May 2003, data would be collected in the directory named $TANKDIR/pmf/sanfs/20030509_080515.

You can use the **–outputDir** parameter to specify the output directory in which the nested subdirectories are created. Uniquely-named output directories is another way to preserve the results of multiple invocations.

It is recommended that the output directory be a mount point for a separate file system. Using a dedicated file system avoids the possibility of filling critical disk space with the results of multiple invocations of this command. This command does some basic checking to ensure that sufficient space exists for the collected data; however, removing old collections of data is your responsibility.

Clients for Windows

For clients for Windows, the **pmf** command attempts to copy any SAN File System trace files (c:\Program Files\IBM\Storage Tank\client\log\sanfs.log) and system dump files (memory.dmp). The output of this command is saved in the c:\Program Files\IBM\Storage Tank\pmf*computer_name**user_ID*\diag_data file directory, where *computer_name* is the host name of the machine for which data is being collected and *user_ID* identifies the user who invoked the command.

Note: This file is overwritten each time you run the **pmf** command. This command copies additional data to this directory, as well.

Additional steps are required to manually collect additional data that is collected by the **pmf** command on Linux systems and clients for AIX. These steps must be performed after the **pmf** command has run.

1. Collect disk information:
   a. Double-click the My Computer icon.
   b. For each logical disk drive, perform the following steps:
      1) Record the drive letter.
      2) Click the drive letter.

3) Select **File → Properties**.
4) From the General tab, record the following information:
   - Label
   - Type
   - File system
   - Used space
   - Free space

2. Collect operating system information:
   a. Select **Start → Settings → Control Panel**.
   b. Double-click **System** to display the System Properties window.
   c. Under the General tab, record the operating system information (for example, Microsoft Windows 2000, 5.00.2195 Service Pack 3).
3. Collect the event logs:
   a. Select **Start → Settings → Control Panel**.
   b. Double-click **Administrative Tools**.
   c. Double-click **Event Viewer**.
   d. Click **Application Log** to highlight it.
   e. Select **Action → Save Log File As**.
   f. In the **File name** field, type **c:\Program Files\IBM\Storage Tank\pmf***computer_name**user_ID***diag_data*_file_name_**.evt**, where *computer_name* is the host name of the machine for which data is being collected, *user_ID* identifies the user that invoked the command, and ***file_name*** is the name of the file.
   g. Repeat steps 3e and 3f and save the files as a text file (select **Text** in the **Save as type** drop-down list.)
   h. From the Event Viewer window, Click **System Log** to highlight it.
   i. Repeat steps 3e through 3g to collect the System log.
4. Collect the SAN File System configuration information:
   a. Select **Start → Run**.
   b. Type **regedit** and click **OK**.
   c. Select **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\stfs\Parameters** directory.
   d. Copy the values for the following parameters:
      - ClientName
      - InstanceName
      - PreferredDriveLetter
      - ServerIpAddress
      - ServerIpPort

   Note: You can export the entire registry so that all information is available. Because the registry might contain customer-sensitive data, so you should get permission from the customer before hand.

Privileges

The **pmf** command relies upon a number of SAN File System and operating system commands that require certain administrative privileges to be executed. To achieve the full benefit of the command, you must be logged in as root on a UNIX platform or Administrator on a Windows platform.

Implementation

The **pmf** command is supplied as a collection of scripts. It is invoked as a shell script (pmf.sh) on UNIX and as a batch file (pmf.bat) on Windows. For Linux systems and clients for AIX, the interface and basic function of the command are equivalent. The interfaces for these platforms invoke an underlying Perl program (pmf.pl), which in turn relies on a collection of additional Perl files for its execution. Therefore, Perl must be installed locally on those machine on which the command is to be used.

For clients for Windows, the pmf command is a batch utility and does not require Perl.

Collected data

The **pmf** command collects data for Metadata servers, Administrative servers, and clients of a SAN File System configuration, regardless of platform. The command relies both on SAN File System administrative commands and on operating-system commands to obtain much of the data it collects. It creates an ASCII file in the time-stamped output directory to record the results of the administrative and operating system commands that it invokes. In the event that the process from which data is requested is unavailable (for example, if the Metadata server is unavailable or unresponsive), the command records a suitable message and continues without collecting data from the unavailable process. This command also assembles various configuration, log, trace, and core files in the output directory.

The **pmf** command collects information related to both hardware and software. The specific information gathered depends on the platform, the SAN File System configuration of the machine, and the parameters specified for that execution. In general, this command collects more data for Metadata servers than for clients, due to the greater function and complexity of the Metadata server configuration.

For hardware, this command collects information about:

- Processors
- Memory
- Network adapters (LAN and SAN)
- Storage devices (local and SAN)

For software, this command collects information about:

- Operating system
- Network configuration (LAN and SAN)
- Storage configuration (local and SAN)
- SAN File System configuration

Within these general categories, the command gathers essential component information (for example, name, version, and configuration), current status, logged messages, trace and diagnostic output, core files, and performance statistics and metrics, among other things. Not all types of information pertain to all hardware and software components about which information is collected.

Examples

Collecting default, system core, and binary data The following example collects the default data, a system core file, and the binary file for the operating system

that generated it, from the local client. The collected data is archived and compressed into a single output file for delivery to remote support personnel.

```
/usr/tank/client/bin/pmf -collect "full,osdump" -archive
```

Related topics

- "collectdiag" on page 181

startCimom

Starts the Administrative agent.

Syntax

Description

Prerequisites:

1. You must have root privileges to use this command.
2. This task must be performed only by trained service technicians.

Note: This command is run from the shell prompt. It is not run inside of tanktool.

This utility assumes that the required files are located in the /usr/tank/admin directory.

Examples

Start the Administrative agent The following example starts the Administrative agent.

```
/usr/tank/admin/bin/startCimom
```

Related topics

- "stopCimom" on page 351

startConsole

Starts the SAN File System console Web server.

Prerequisite: You must have root privileges to use the command.

Syntax

Description

Note: This command is run from the shell prompt. It is not run inside of tanktool.

This command is case sensitive and must be entered as shown.

The Web server must be running for you to open the SAN File System console. The Web server starts automatically when the storage engine boots up. This command allows you to manually restart the Web server if it is stopped.

Examples

Starts the Web server The following example starts the SAN File System console Web server.

```
startConsole
Starting the SAN File System console...
The SAN File System console is operational at https://189.24.15.162:7979/tank
```

Related topics

- "stopConsole" on page 352

stfsdebug

Enables or disables the logging of file-system-driver debug messages in the syslog for the specified virtual client for AIX.

Syntax

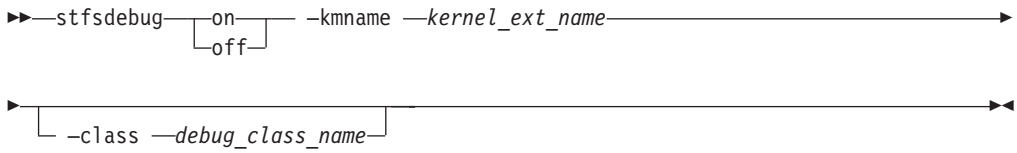

Parameters

on Enables the logging debug messages in the syslog.

off
Disables the logging of debug messages in the syslog.

–kmname *kernel_ext_name*
Identifies kernel-extension name of the file-system driver associated with the virtual client.

The file-system driver is loaded as a kernel extension. To identify the instance of the file-system drive, you identify the kernel extension. Each kernel extension has a name, but this name is not unique. This name is usually the file name of the object file from which you loaded the kernel extension (for example, /usr/tank/client/bin/stfs).

–class *debug_class_name*
Identifies the class name for debug messages. If not specified, this command applies all classes. You can specify any of the following values:

ATTR File attribute information.

CONFIG
Driver initialization, termination, and control.

DISK Disk configuration at client-creation time

IO I/O operations.

MOUNT
FSIs and clients creation and destruction.

PAGER
Pager and items that give work to the pager.

RDWR
Reading from and writing to file cache.

RNGLOCK
Range lock information.

VFSOP
vfs operations, excl get root vnode

VNODE
Management of vnodes, gnodes, and stfsnodes.

VNODEOP
vnode operations, excl lookup, hold and release vnode.

Description

Prerequisite: This task must be performed only by trained service technicians.

Examples

Enable debugging on a client for AIX The following example enables the logging of file-system-driver debug messages in the syslog for the specified virtual client.

```
stfsdebug on –kmname /lib/stfs
```

Related topics

- "Clients" on page 10
- "stfsdriver" on page 309
- "UNIX-based clients" on page 12

stfsstat

Displays statistics for the local client.

Syntax

Clients for AIX

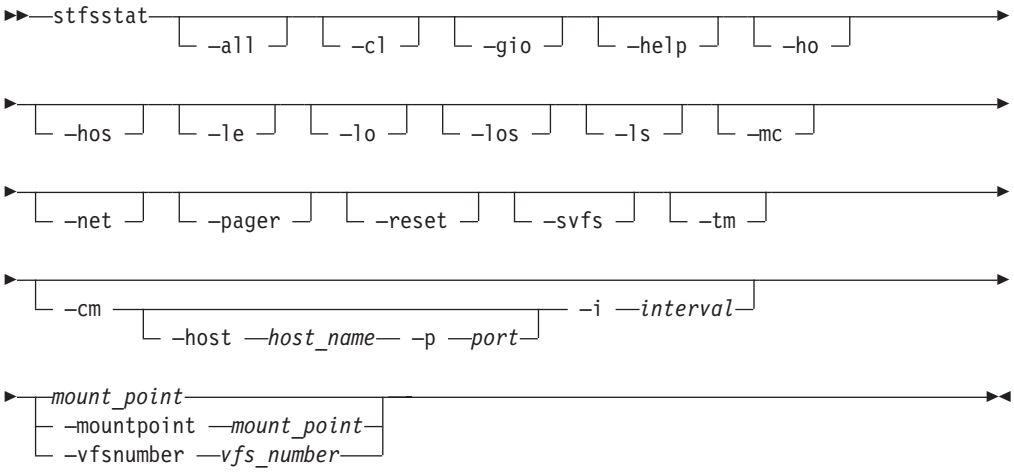

Clients for Windows

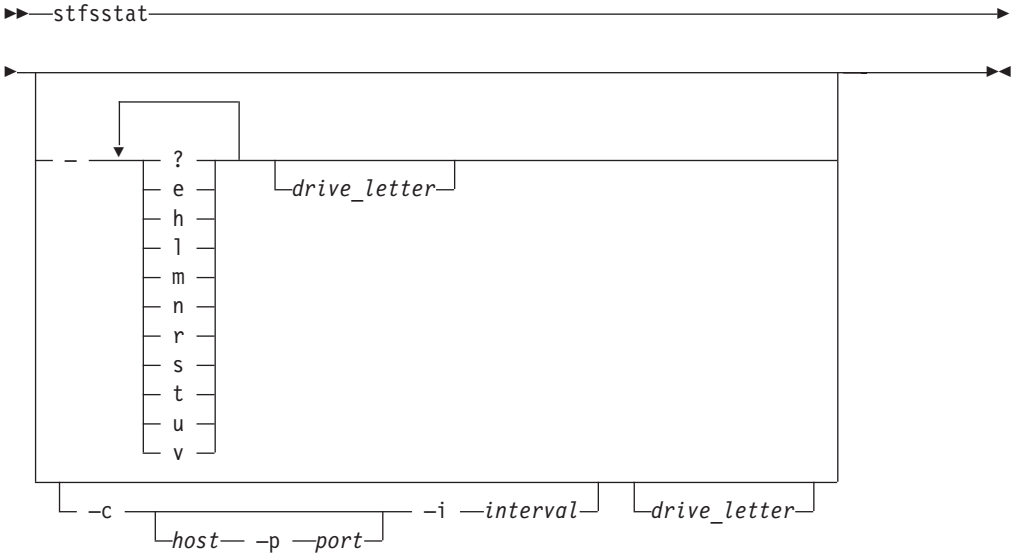

Parameters

Clients for AIX

–all
Displays all client for AIX statistics, except per-object statistics. This parameter is the same as specifying **–tm**, **–mc**, **–ls**, **–le**, **–svfs**, **–gio**, **–pager**, and **–cl** parameters.

–cl
Displays cleaner statistics.

–cm [–host ***host_name*** **–p** ***port***] **–i** *interval*
Continuously monitors the specified statistics. The statistics are sent as a string

to the specified *host* on the specified UDP ***port*** at every specified *interval* (in seconds). The minimum interval is 20 seconds. If no host or port is specified, the statistics are sent to stdout.

–gio
Displays I/O statistics per global disk (GDISK), including : blocks read and written, and the average number of bytes read and written.

–help
Displays a detailed description for this command.

–ho
Lists the states of all objects in the mcObject hash table, including object ID, flags, and attributes

–hos
Displays a short list of object hash table.

–le
Lists mutex and latch state of all objects in the mcObject hash table, including current holder, hold state, and waiters.

–lo
Lists the states of all objects in the mcObject LRU list, including object ID, flags, and attributes.

–los
Displays a short list of object LRU.

–ls
Lists mutex and latch statistics of metadata cache objects such as mcInstance and mcObject. Statistics include the number of attempts, wait time, and hold time.

–mc
Displays metadata-cache statistics.

–net
Displays network statistics per Metadata server with which this client has a lease.

–pager
Displays pager statistics per client, including information about the page-in and page-out activity.

–reset
Resets all statistics to zero. Some statistics, such as object and latch state, cannot be reset because no counters are incremented, but data structures are shown.

–svfs
Displays file-system statistics. These metrics are mount-specific and distinguish between the different mount points. Statistics include the number of file, inode, address space, and super operations performed.

–tm
Displays Transaction Manager (TM) statistics per client related to TM data structures, including the number of messages sent and received (per message type), the maximum lengths and average lengths of the transaction queue, and the number of transactions, messages, and leases lost.

mount_point | **–mountpoint** *mount_point*
The mount point where you want to mount the file-system image (for example, the directory above which you want the global namespace image's directory tree to appear).

–vfsnumber *vfs_number*
Identifies the virtual-file-system (VFS) number associated with the global namespace image for which you want to display statistics. In AIX, every global namespace image has a unique VFS number.

The **stfsmount** command displays this number when it creates the global namespace image.

Clients for Windows

–? Displays a detailed description for this command.

–c Continuously monitors the specified statistics.

–e Displays the latch state.

–h Displays the object hashlist statistics.

–i Displays the interval for sending statistics. This option must be used with the **–c** option.

–l Lists mutex and latch statistics of metadata cache objects.

–m
Displays metadata-cache statistics.

–n Displays the engine statistics.

–p Displays the port for sending statistics. This option must be used with the **–c** option.

–r Resets the CSM statistics

–s Displays the object LRU statistics.

–t Displays Transaction Manager (TM) statistics. This is the default option.

–u Displays statistics in the short-output format.

–v Displays file-system statistics

drive_letter
The drive letter mapped to SAN File System. The default letter is T.

Description

Prerequisites: This task must be performed only by trained service technicians.

On Windows machines running the client, the statistics options (**?**, **e**, **h**, **l**, **m**, **n**, **r**, **s**, **t**, **u**, and **v**) are appended together. For example, to choose to display the latch state and statistics, you would specify **–el**.

Examples

Clients for AIX

Displays statistics The following example displays transaction-manager statistics and metadata-cache statistics for the local client that is mounted on /mnt/sanfs.

```
stfsstat –tm –mc /mnt/sanfs
```

Clients for Windows

Displays statistics The following example displays statistics for the object-hash list and latch state the local client mapped to drive S.

```
stfsstat –he s
```

Related topics

- "Clients" on page 10
- "stfsclient" on page 303
- "stfsdriver" on page 309
- "UNIX-based clients" on page 12

stopCimom

Stops the Administrative agent.

Syntax

Description

Prerequisites:

1. You must have root privileges to use this command.
2. This task must be performed only by trained service technicians.

Note: This command is run from the shell prompt. It is not run inside of tanktool.

This utility assumes that the required files are located in the /usr/tank/admin directory.

Examples

Stop the Administrative agent. The following example stops the Administrative agent.

```
/usr/tank/admin/bin/stopCimom
```

Related topics

- "startCimom" on page 343

stopConsole

Stops the SAN File System console Web server.

Syntax

Description

Prerequisite: You must have root privileges to use the command.

Note: This command is run from the shell prompt. It is not run inside of tanktool.

This command is case sensitive and must be entered as shown.

You would stop the SAN File System console Web server when you are upgrading software.

Examples

Stops the Web server The following example stops the SAN File System console Web server.

```
stopConsole
Stopping the SAN File System console...
```

Related topics

- "startConsole" on page 344

tank extractbootrecord

Extracts the product (disk) label information contained in the master volume and creates a local copy of the Tank.Bootstrap file.

Syntax

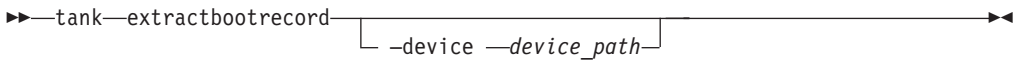

Parameters

–device *device_path*
Specifies the device path for a valid master volume from to read the product label.

Description

Prerequisites::

1. This task must be performed only by trained service technicians.
2. The cluster must be offline.

Note: This command is run from the shell prompt. It is not run inside of tanktool.

This command allows a single-engine cluster to be brought up on new hardware in the case of disaster recovery or when you attach new hardware to an existing cluster. Recovering the Tank.Bootstrap file is required only when copies on all engine hosting Metadata servers in the cluster are lost.

This command is located in the /usr/tank/server/bin directory.

This command is used in disaster recovery situations or when attaching new hardware to an existing set of SAN devices.

Each Metadata server in the cluster stores a local copy of the product label that is on the master volume. The local copy is stored in the file /etc/tank/server/Tank.Bootstrap, in binary format. SAN File System uses this information to locate the master volume on the SAN at boot time. After it is located, the cluster definition can be read. You would use this command when the local copy of the Tank.Bootstrap file is lost on all engines to extract the product label from the system master volume and recreate the local Tank.Bootstrap file.

This command is used to recreate the local Tank.Bootstrap file on the master Metadata server in a single-engine cluster. The Tank.Bootstrap file is created on each remaining Metadata server as they are added to the cluster (using the **addserver** command). It is not necessary or recommended to run this command on every Metadata server or to copy this file from one Metadata server to another.

Examples

Extract a boot record The following example extracts a local Tank.Bootstrap file from device /dev/rsdc. This device is a valid master volume of a SAN File System cluster instance.

```
#/usr/tank/server/tank extractbootrecord –device /dev/rsdc
Label information from master disk /dev/srdc was extracted and stored in
Tank.Bootstrap.
```

Related topics

- "addserver" on page 322
- "dropserver" on page 326
- "tank lsdisklabel" on page 358
- "tank resetcluster" on page 360
- "startserver" on page 276
- "stopserver" on page 293

tank lscluster

Displays the cluster definition from a system master volume when the Metadata servers are not running.

Syntax

Description

Prerequisites:

1. This task must be performed only by trained service technicians.
2. The cluster must be offline.

Note: This command is run from the shell prompt. It is not run inside of tanktool.

This command is used in disaster recovery situations, while troubleshooting the system, or to inspect the cluster definition.

This command is located in the /usr/tank/server/bin directory.

This command displays the following information:

- Master state of the state machine in the cluster services component.
- Subordinate state of the state machine in the cluster services component.
- Transport protocol used for group services. The value defaults to UDP and cannot be changed. Note that the group services includes cluster services.
- Transition state, which indicates whether a cluster transition is in progress at the time this command is issued.
- Started state, which indicates whether the group services subsystem was started.
- Joining state, which indicates whether the Metadata server attempted to join the cluster.
- Forming state, which indicates whether the Metadata server is attempting to reform the cluster as the master Metadata server.
- Dynamic group ID, which identifies the subset of the cluster that is alive and well.
- Time when the cluster was last reformed (committed).
- Size of the current dynamic group. This is affected when the cluster or a Metadata server starts, stops, crashes, or aborts.
- Size of the entire static cluster. This is affected by the **addserver** and **dropserver** commands.
- Cluster identifier, which is initialized at installation time.
- Cluster name (sanfs).
- Time of the installation.
- Timeout value used by group services.
- IP address of the network used by the cluster.
- Netmask used by the cluster.
- Current active version of the cluster boot record. Two copies are maintained. This value will be A or B.
- Identifier of the engine hosting the Metadata server where this command was issued.

- A list of all engines in the cluster and attributes for each. The attributes include:
  - Engine identifier.
  - IP address.
  - Group services port number (gs)
  - SAN File System protocol port number (stp).
  - Heartbeat port number (hb).
  - Administrative port number (adm)
  - Administrative agent port number (agent)
  - Metadata server/engine name.
  - Last committed software version.

Examples

Lists the static cluster definition The following example lists the static cluster definition before running the **tank resetcluster** command.

```
tank lscluster
Group information:
mast_state:          Microkernel
sub_state:           Invalid
protocol:            udp
in transition:       no
b_started:           0
b_joining:           0
b_forming:           0
group_id:            9
group commit time:   Jun 25, 2003 9:47:53 PM
group size:          4
cluster size:        4
cluster id:          1234
cluster name:        sanfs
install time:        Jun 25, 2003 4:23:45 PM
installation id:     835407414733488113
gs_timeout:          1000 (millisecs)
ip network:          0.0.0.0 (interfaces: )
netmask:             0.0.0.0
active set:          B
this node id:        0

      id   ip addr         gs      stp     hb      gdm     agent
NODE: 0    192.168.10.88   11003   11001   11004   11002   5989
NODE: 1    192.168.11.88   11003   11001   11004   11002   5989

stnode name   SW Version
GR ST0        1.0.1
GR ST1        1.0.1
```

Related topics

- “addserver” on page 322
- “dropserver” on page 326
- “tank extractbootrecord” on page 353
- “lsserver” on page 221
- “tank resetcluster” on page 360
- “statserver” on page 286
- “statcluster” on page 278

tank lsversion

Displays the version control information from a system master disk.

Syntax

Description

Prerequisites::

1. This task must be performed only by trained service technicians.
2. This command must be used only under the direction of your IBM support representative.

You would use this command only during troubleshooting or disaster recovery situations.

Examples

Displays version control information The following example version control information from a system master disk.

```
/usr/tank/server/bin/tank lsversion
```

Related topics

- "tank resetversion" on page 362

tank lsdisklabel

Displays the a SAN File System product (disk) label for a specified device.

Syntax

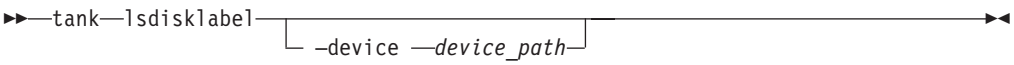

Parameters

–device *device_path*
Specifies the device path of a *valid* master volume from which to read a product label. If not specified, the product label that is read from the local /etc/tank/server/Tank.Bootstrap file, which contains a subset of the product label found on the master volume.

Description

Prerequisites::

1. This task must be performed only by trained service technicians.
2. The cluster must be offline.

Note: This command is run from the shell prompt. It is not run inside of tanktool.

This command is used in disaster recovery situations, when you attach new hardware to an existing cluster, or while troubleshooting the system.

This command is located in the /usr/tank/server/bin directory.

If you display the product label from the master volume, this command displays this information:

- Label identifier 1. This is always SDISK for all SAN File System devices. This should match the label identifier 2. It is used as an integrity check to detect valid or corrupted product labels.
- Label number.
- Disk type. Valid values are:
  M Master volume label
  S System volume label
  U User volume
- Global disk identifier. This is the ID of volume that is used and understood by the Metadata servers and clients.
- Label date. This is the date when the volume was added.
- Owner identifier. This is the ID of the cluster that owns the volume.
- Installation identifier. This is the unique installation ID for the shared storage that is initialized during installation and changes only upon reinstallation.
- Disk epoch identifier.
- Product identifier.
- Label identifier 2. This is always SDISK for all SAN File System devices. This should match the label identifier 1. It is used as an integrity check to detect valid or corrupted product labels.

If you display the product label from a local Tank.Bootstrap file, this command displays this information:

- Disk type.
- Global disk identifier.
- Owner identifier.
- Installation identifier.
- Disk epoch identifier.

Note: The owner ID, installation ID, and disk epoch ID for a given volume is the same as the owner ID, installation ID, and disk epoch ID contained in the Tank.Bootstrap file on all engines in the cluster that own that volume. An engine may be zoned in such a way that it sees devices from another cluster, but each cluster uses the product label information in the Tank.Bootstrap file to know which devices it owns.

Examples

Displays the product label for the system master volume The following example displays the product label for the system master volume contained on device /dev/rsdc.

```
/usr/tank/server/bin/tank lsdisklabel –device /dev/rsdc
-------------------------------------------------------
Product Label on Device /dev/rsdx:
-------------------------------------------------------

Label ID 1                  : SDISK
Label Number                : 0001
Disk Type                   : M
Global Disk ID              : 73EFA22D6A2355F1
Label Date                  : Jun 25, 2003 4:35:24 AM
Owner ID                    : 00000000000004D2
Installation ID             : 73EFA22D6A2355F1
Disk Epoch                  : 0000000000000000
Product ID                  : STORAGETANK
Label ID 2                  : SDISK
-------------------------------------------------------
```

Displays the information in Tank.Bootstrap The following example displays the product label information found in the local Tank.Bootstrap file.

```
/usr/tank/server/bin/tank lsdisklabel
-------------------------------------------------------
Tank.Bootstrap information
-------------------------------------------------------

Disk Type                   : U
Global Disk ID              : 73F02ABDBE27B00B
Owner ID                    : 00000000000004D2
Installation ID             : 73F02ABDBE27B00B
Disk Epoch                  : 0000000000000000
-------------------------------------------------------
```

Related topics

- “addserver” on page 322
- “tank extractbootrecord” on page 353
- “tank resetcluster” on page 360

tank resetcluster

Erases the static cluster definition contained on the system master volume, without reinitializing the metadata, which in affect drops all Metadata servers from the cluster at one time.

Syntax

Description

Prerequisites:

1. This task must be performed only by trained service technicians.
2. The cluster must be offline.

Note: This command is run from the shell prompt. It is not run inside of tanktool.

Attention: This command will wipe out the static cluster definition contained on the master volume. It is equivalent to issuing the **dropserver** command on all engines.

This command is used in extreme disaster recovery situations or when you attach new hardware to an existing cluster.

This command is located in the /usr/tank/server/bin directory.

This command must be run from an engine that has a Tank.Bootstrap file. If the engine does not have a Tank.Bootstrap file, this file can be recovered from the master volume using the **tank extractbootrecord** command. If the master volume is unknown, you can issue the **tank lsdisklabel** command to inspect the SAN File System product label and to locate the master volume.

After you run this command, the first Metadata server that is started with a valid Tank.Bootstrap file becomes the new master Metadata server. At this point, the cluster is a single-engine cluster, and all additional Metadata servers must be added using the **addserver** command. In addition, it might be necessary to move filesets to these new engines (using the **setcontainerserver** command) if the added engines have different names than the original cluster's engine names.

Examples

Reset the cluster The following example erases the static cluster definition on the master volume.

```
#/usr/tank/server/bin/tank resetcluster
Static cluster definition has been reset. cluster size = 0
```

Related topics

- "addserver" on page 322
- "dropserver" on page 326
- "tank extractbootrecord" on page 353
- "tank lscluster" on page 355
- "tank lsdisklabel" on page 358
- "setcontainerserver" on page 260
- "startserver" on page 276

- “stopserver” on page 293

tank resetversion

Resets the version-control information on a system master disk.

Syntax

Description

Prerequisites::

1. This task must be performed only by trained service technicians.
2. The cluster must be offline.

Attention: This command should be used only by trained service technicians and only in the event of extreme disaster recovery or hardware refresh situations. Invoking this command at runtime can lead to unpredictable and destructive results.

Examples

Resets the version-control information The following example resets the version-control information on a system master disk.

```
/usr/tank/server/bin/tank resetversion
```

Related topics

- "tank lsversion" on page 357

Appendix D. Engines

This topic describes the components of the SAN File System engine.

Major components of the Model 1RX

The serial number, machine type number, and model number for the Model 1RX are located on the label that is attached to the bezel on the front of the machine. You will need these numbers when you register your machine with IBM. The information label containing the machine type number, model number, and agency mark for the Model 1RX are located on the bottom of the machine.

The orange color on components and labels in the Model 1RX identifies hot-swap or hot-plug components. You can install or remove hot-swap or hot-plug components while the system is running, provided that your system is configured to support this function.

The blue color on components and labels indicates touch points where a component can be gripped, a latch moved, and so on.

The following illustration shows the major components in the Model 1RX.

Note: The illustrations in this document might differ slightly from your hardware.

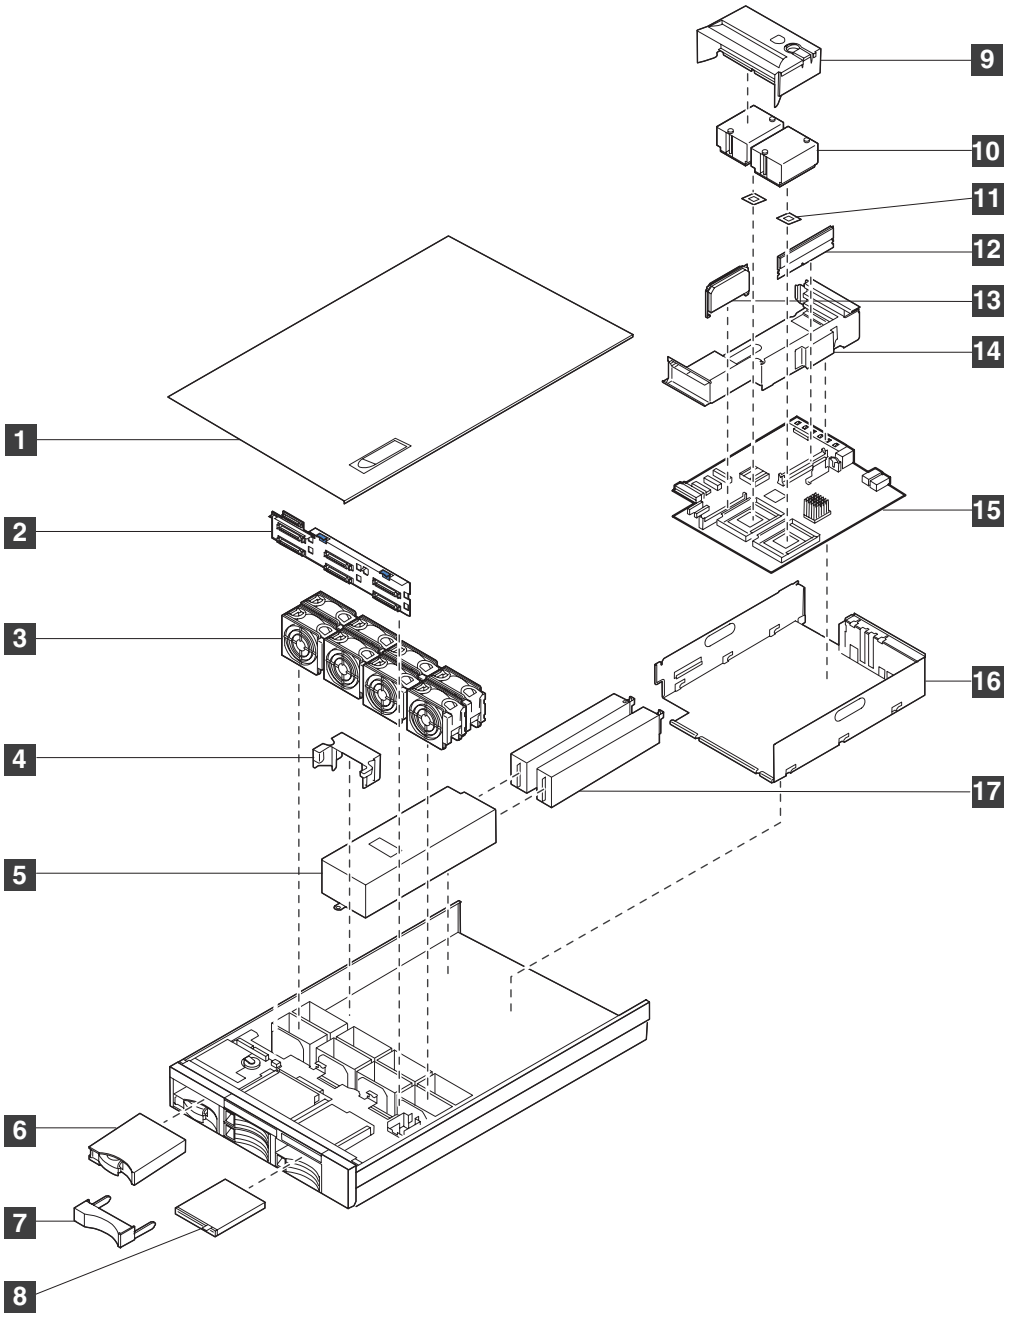

1 Top cover
2 DASD backplane with bracket
3 Fan, hot-swap
4 Thermal baffle (part of kit with item 9)
5 Power cage and backplane assembly
6 Hot-swap hard disk drive
7 Filler panel for drive bay
8 CD-ROM
9 Thermal baffle (part of kit with item 4)
10 Microprocessor air baffle and Heat sink
11 Microprocessor
12 Memory module
13 VRM
14 PCI riser card assembly

15 System board
16 Shuttle assembly
17 Hot-swap power supply

Related topics:

- "Controls and indicators"
- "System-planar external port connectors" on page 367
- "System-planar internal cable connectors" on page 368
- "System-planar LED locations" on page 369
- "System-planar option connectors" on page 371
- "System-planar switches and jumpers" on page 369
- "PCI riser-card option connectors" on page 372
- "Replacing storage engine components" on page 411
- "Replacing master console components" on page 411
- Chapter 11, "Adding and replacing engine components", on page 109

Controls and indicators

On the Model 1RX, the controls and indicators are located on the front and the rear of the engine.

Related topics:

- "Model 1RX front view"
- "Model 1RX rear view" on page 366

Model 1RX front view

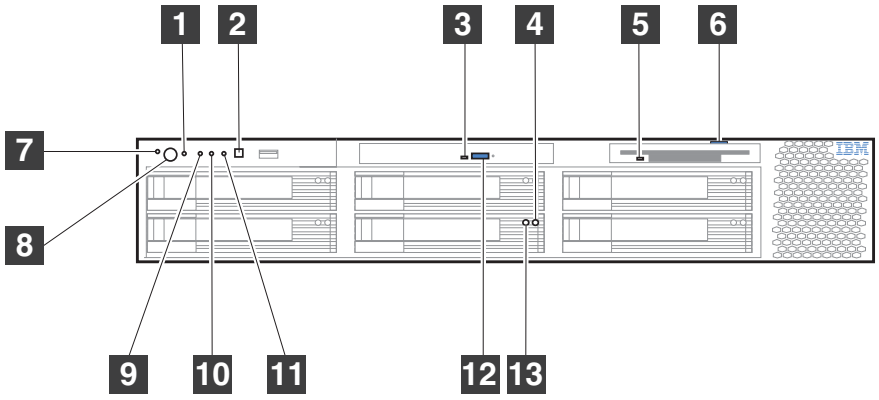

1 **Reset button:** Press this button to reset the engine and run the power-on self-test (POST). You might need to use a pen or the end of a straightened paper clip to press the button.

2 **System-error LED:** This amber LED is lit when a system error occurs. An LED on the diagnostic LED panel will also be on to further isolate the error.

3 **CD-ROM drive activity LED:** This LED is lit when the CD-ROM drive is in use.

4 **Hard disk drive status LED:** Each hot-swap drive has a hard disk drive status LED. This amber LED is lit continuously when the drive has failed.

5 **Diskette drive activity LED:** This LED is lit when the diskette drive is in use.

6 **Diskette-eject button:** Press this button to release a diskette from the drive.

7 **Power-on LED:** This green LED is lit and stays on when you power on the engine and flashes when the engine is in Standby mode.

Attention: If the power-on light-emitting diode (LED) is off, it does not mean there is no electrical current present in the engine. The LED might be burned out. To remove all electrical current from the engine, you must unplug the power cords from the electrical outlets or from the uninterruptible power device.

8 **Power-control button:** Press this button to manually turn on the engine and put the engine in Standby mode.

9 **SCSI or IDE bus activity LED:** This LED is on when there is activity on the SCSI or IDE bus.

10 **System-identification LED:** This blue LED is lit in response to a programmed condition, or it can be turned on remotely by the system administrator to aid in identifying a specific engine for maintenance.

11 **Information LED:** This amber LED is lit when the information log contains information about certain conditions in the engine that might affect performance.

12 **CD-eject button:** Press this button to release a CD from the drive.

13 **Hard disk drive activity LED:** Each of the hot-swap drives has a hard disk drive activity LED. This green LED flashes when the controller is accessing the drive.

Related topics:

- "Input and output ports" on page 372
- "Major components of the Model 1RX" on page 363
- "Model 1RX rear view"

Model 1RX rear view

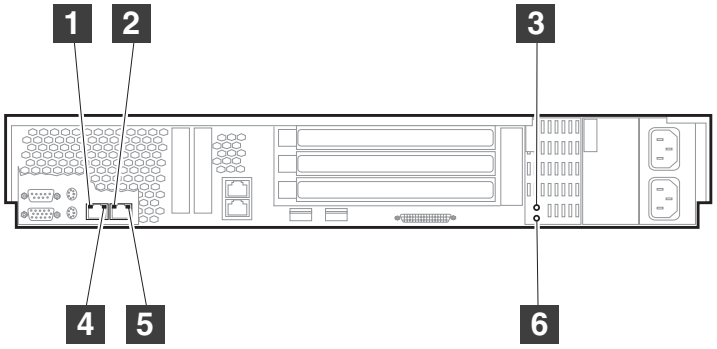

1 **Ethernet 2 activity LED:** This green LED is lit when the engine is transmitting or receiving signals to the Ethernet LAN that is connected to Ethernet port 2.

2 **Ethernet 1 activity LED:** This green LED is lit when the engine is transmitting or receiving signals to the Ethernet LAN that is connected to Ethernet port 1.

3 **AC power LED:** This green LED provides status information about the power supply. During typical operation, both the ac and dc power LEDs are lit.

4 **Ethernet 2 link status LED:** This green LED is lit when there is an active link connection on the 10BASE-T, 100BASE-TX, or 1000BASE-TX interface for Ethernet port 2.

5 **Ethernet 1 link status LED:** This green LED is lit when there is an active link connection on the 10BASE-T, 100BASE-TX, or 1000BASE-TX interface for Ethernet port 1.

6 **DC power LED:** This green LED provides status information about the power supply. During typical operation, both the ac and dc power LEDs are lit.

Related topics:

- "Input and output ports" on page 372
- "Model 1RX front view" on page 365
- "Major components of the Model 1RX" on page 363

System board

This topic describes the system board for the SAN File System engine.

System-planar external port connectors

The following illustration shows the location of the external input and output port connectors on the system planar.

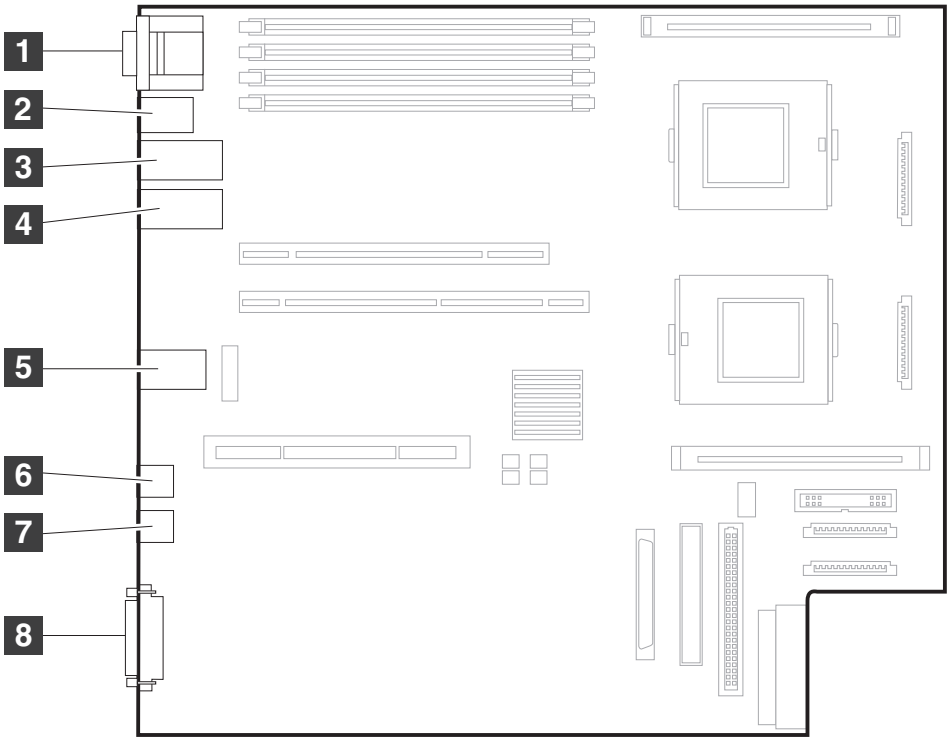

1 Serial and video (J9)
2 Mouse and keyboard (J10)
3 Ethernet 1 (J11)
4 Ethernet 2 (J12)
5 RS485–A and B (J15)
6 USB 1 (J17)
7 USB 2 (J18)
8 External SCSI (J19)

System-planar internal cable connectors

The following illustration shows the internal cable connectors on the system planar.

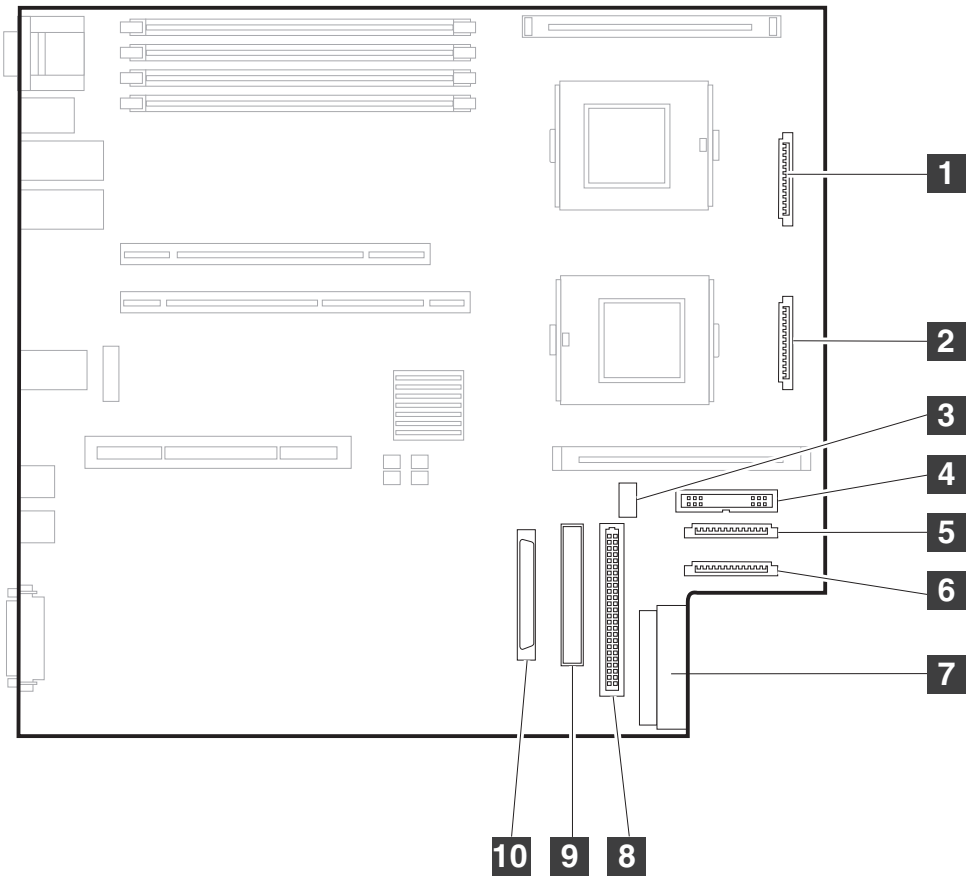

1 Fans 4 and 8 (J28)
2 Fans 3 and 7 (J27)
3 USB 3 operator panel
4 Power signal (J25)
5 Fans 2 and 6 (J29)
6 Fans 1 and 5 (J26)
7 Power (J24)
8 Operator panel (J22)
9 IDE (J21)
10 Internal SCSI (J23)

System-planar LED locations

The following illustration shows the location of the LEDs on the system planar. You might need to refer to this illustration when solving problems with the engine.

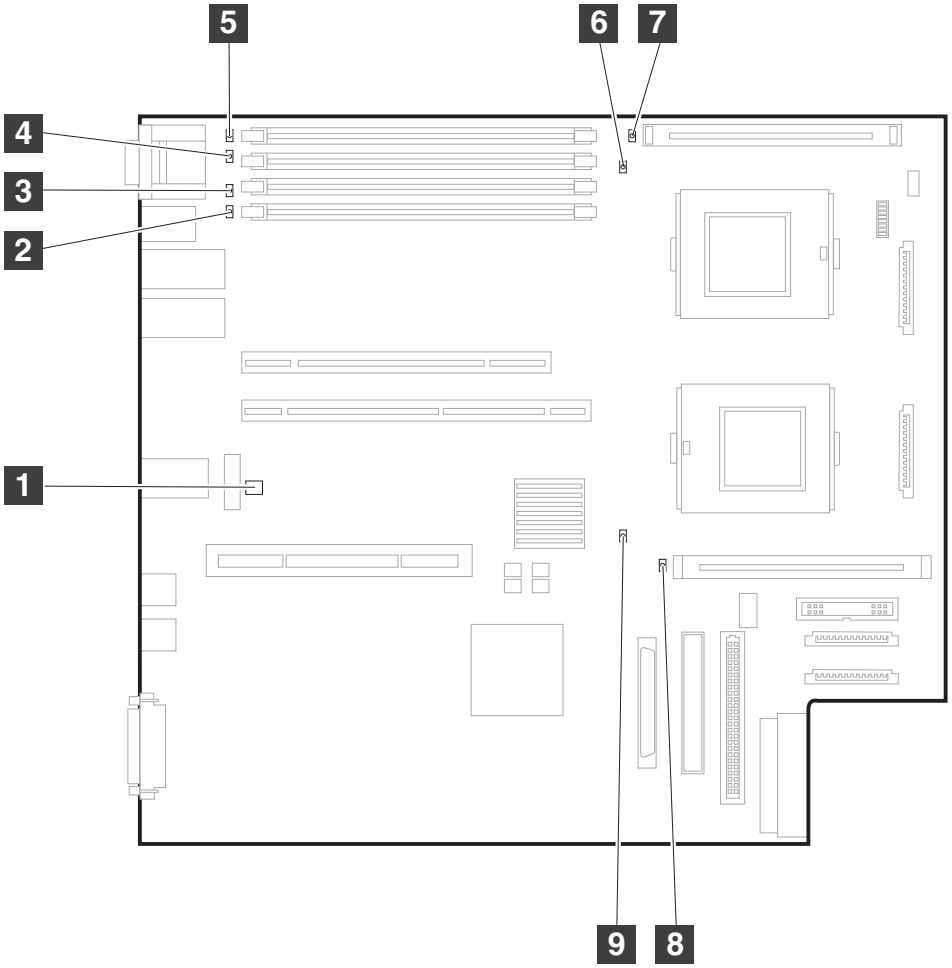

1. DIMM 1 error LED (CR12)
2. DIMM 2 error LED (CR15)
3. DIMM 3 error LED (CR14)
4. DIMM 4 error LED (CR13)
5. VRM 2 error LED (CR16)
6. Microprocessor 2 error LED (CR17)
7. Microprocessor 1 error LED (CR18)
8. VRM 1 error LED (CR19)
9. Integrated System Management Processor (ISMP) activity LED (CR41)

System-planar switches and jumpers

The following illustration shows the location of the switch block and the flash boot jumper on the system planar.

Notes:

1. Before changing any switch settings or moving any jumpers, power OFF the system and any peripheral devices, and disconnect all external cables and power cords.
2. Any system-planar switch or jumper blocks that are not shown in the illustrations in this book are reserved.

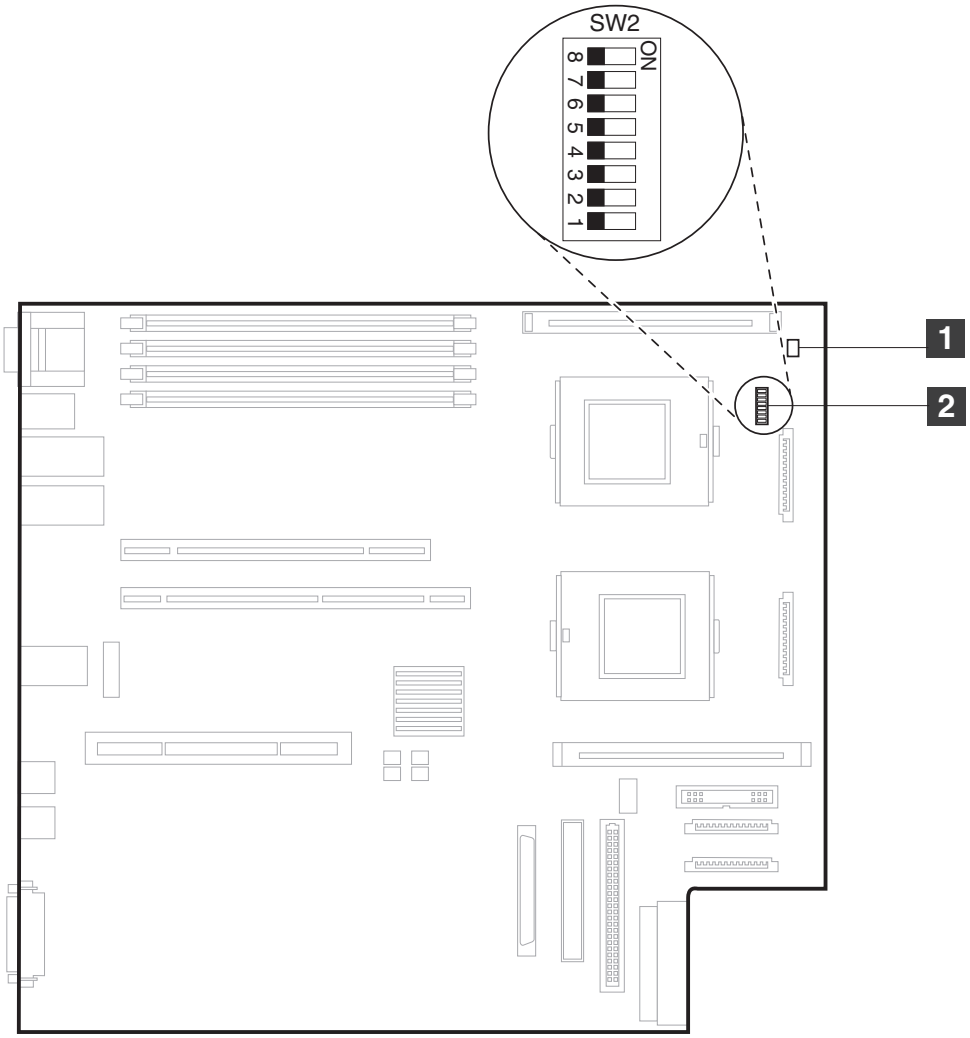

1 System-planar switch block (SW2)
2 Flash-boot-block recovery jumper (J34)

Table 8 describes the function of each switch on the switch block.

Table 8. Switches 1-8

Switch number	Default value	Switch description
8	Off	Reserved.
7	Off	Reserved.
6	Off	Reserved.
5	Off	Reserved.
4	Off	Reserved.
3	Off	Power-on override. When toggled to On, this switch forces the power on, overriding the Power-on button.

Table 8. Switches 1-8 (continued)

Switch number	Default value	Switch description
2	Off	Power-on password override. Changing the position of this switch bypasses the power-on password check the next time the engine is powered on and starts the Configuration/Setup Utility program so that you can change or delete the power-on password. You do not need to move the switch back to the default position after the password is overridden. Changing the position of this switch does not affect the administrator password check if an administrator password is set.
1	Off	H8 out serial port. This switch is reserved for service use only.

Related topics:

- "Safety information" on page 418
- "Starting the Configuration/Setup Utility" on page 106

System-planar option connectors

The following illustration shows the location of the system-planar connectors for user-installable options.

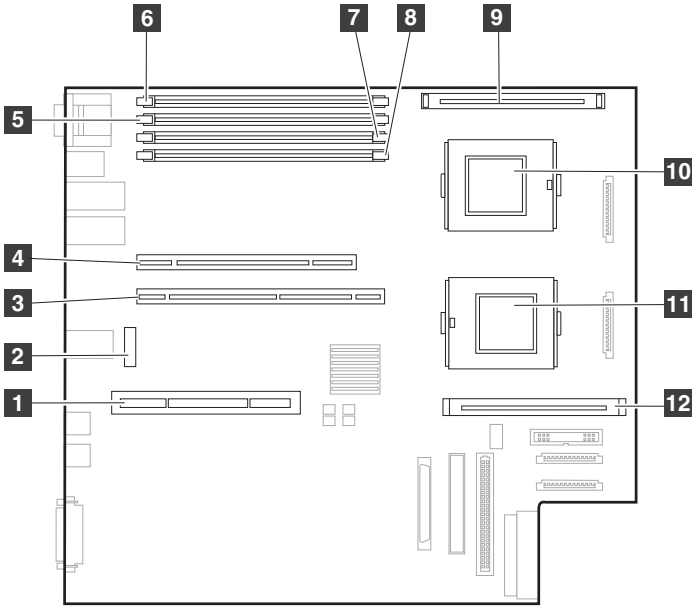

1 PCI riser card (J16)
2 Battery (BH1)
3 PCI-X slot 2 64-bit 3.3V 100 MHz (J13)
4 PCI-X slot 1 64-bit 3.3V 100 MHz (J14)
5 DIMM 2 (J2)
6 DIMM 1 (J1)
7 DIMM 3 (J3)
8 DIMM 4 (J4)
9 VRM 2 (J5)
10 Microprocessor 2 (J6)
11 Microprocessor 1 (J7)
12 VRM 1 (J8)

PCI riser-card option connectors

The following illustration shows the location of the PCI riser-card connectors for user-installable options.

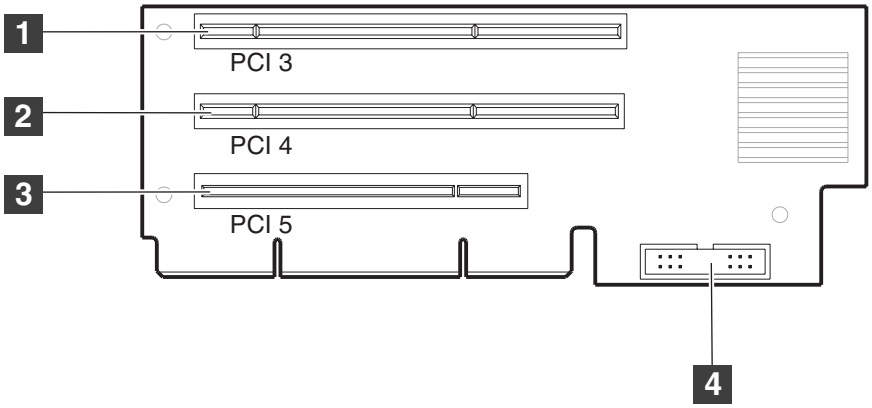

1. PCI-X slot 3 64-bit 3.3V 133 MHz (PCI 3)
2. PCI-X slot 4 64-bit 3.3V 133 MHz (PCI 4)
3. PCI-X slot 5 32–bit 5V 33 MHz (PCI 5)
4. System management connector

Input and output ports

The I/O ports are on the rear of the Model 1RX. These ports include the following:

- One video port
- One keyboard port
- One auxiliary-device (pointing device) port
- One dual-channel Ultra320 SCSI controller (LVD) SCSI port
- One serial port
- Three Universal Serial Bus (USB) version 1.1 ports (USB 1 and USB 2 on the rear, USB 3 on the front)
- Two Ethernet ports
- Two RS-485 communication ports dedicated to the Integrated System Management Processor (ISMP)

Rear view

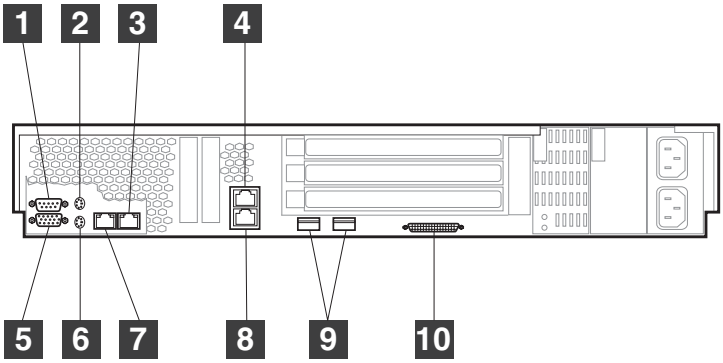

1. Serial port
2. Auxiliary device (pointing device) port
3. Ethernet 2 port
4. RS 485-A (ISMP) port
5. Video port
6. Keyboard port
7. Ethernet 1 port

8 RS 485-B (ISMP) port
9 USB 1 and 2 ports
10 SCSI port

Front view

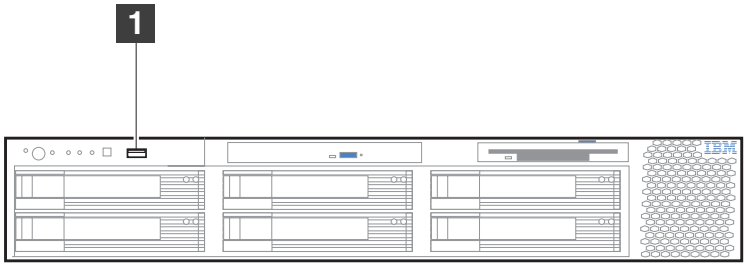

1 USB 3 port

Related topics:

- "Auxiliary-device (pointing device) port"
- "Ethernet port"
- "Integrated system management port" on page 374
- "Keyboard port" on page 374
- "Ultra320 SCSI controller system-planar port" on page 374
- "Universal Serial Bus version 1.1 port" on page 375
- "Video port" on page 375

Auxiliary-device (pointing device) port

The system board has one auxiliary-device port that supports a mouse or other pointing device.

The following illustration shows the auxiliary-device connector on the rear of the Model 1RX. This connector conforms to the industry standard for a PS/2 pointing device.

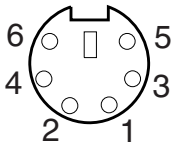

Related topics:

- "Input and output ports" on page 372

Ethernet port

To access the Ethernet controller, connect a Category 5, or higher, unshielded twisted-pair (UTP) cable to either of the Ethernet (RS-45) connectors on the rear of the Model 1RX.

There are four RS-45 Ethernet connectors on the rear of the Model 1RX.

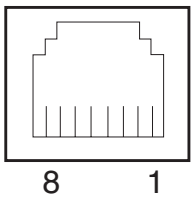

Related topics:

- "Input and output ports" on page 372

Integrated system management port

Related topics:

- "Input and output ports" on page 372

Keyboard port

The following illustration shows the keyboard connector on the rear of the engine. This connector conforms to the industry standard for a PS/2 keyboard.

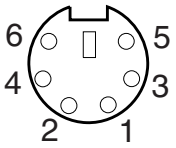

Related topics:

- "Input and output ports" on page 372

Ultra320 SCSI controller system-planar port

The Model 1RX comes with one SCSI cable, which connects the internal connector on the system planar to the standard hot-swap-drive backplane. If you plan to attach external SCSI devices, you must order additional cables. To select and order the correct cables for use with external devices, contact your IBM reseller or IBM marketing representative.

The following illustration shows a 68-pin, female D-shell SCSI connectors. This connector conforms to the SCSI standard.

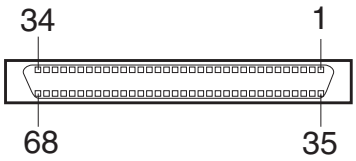

Related topics:

- "Input and output ports" on page 372
- "SCSI IDs"
- "SCSI cabling requirements" on page 375

SCSI IDs:

Table 9. SCSI IDs for standard hot-swap hard disk drives, SCSI controller, and backplane

Device	SCSI ID
Drive bay 1	0
Drive bay 2	1
Drive bay 3	2
Drive bay 4	3
Drive bay 5	4
Drive bay 6	5
SCSI controller	7

Table 9. SCSI IDs for standard hot-swap hard disk drives, SCSI controller, and backplane (continued)

Device	SCSI ID
Backplane	8

The hot-swap-drive backplane controls the SCSI IDs for the internal hot-swap drive bays. However, when you attach an external SCSI device to an optional SCSI adapter, you must set a unique ID for the device. See the information that comes with the device for instructions to set its SCSI ID.

Related topics:

- "SCSI cabling requirements"
- "Ultra320 SCSI controller system-planar port" on page 374

SCSI cabling requirements: Related topics:

- "SCSI IDs" on page 374
- "Ultra320 SCSI controller system-planar port" on page 374

Universal Serial Bus version 1.1 port

USB is a serial interface standard for telephony and multimedia devices. It uses Plug and Play technology to determine the type of device that is attached to the connector.

Notes:

1. If you attach a standard (non-USB) keyboard to the keyboard connector, the USB ports and devices will be disabled during the power-on self-test (POST).
2. If you install a USB keyboard that has a mouse port, the USB keyboard emulates a mouse, and you will not be able to disable the mouse settings in the Configuration/Setup Utility program.
3. Check to make sure that your network operating system (NOS) supports USB devices.
4. For additional information about USB version 1.1 devices, go to http://www.usb.org/.

The following illustration shows a USB-port connector. These connectors conform to the USB version 1.1 standard.

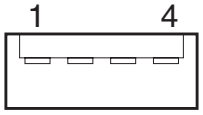

Related topics:

- "Input and output ports" on page 372
- "Starting the Configuration/Setup Utility" on page 106

Video port

Note: If you install a PCI video adapter, the server BIOS will automatically disable the integrated video controller.

The following illustration shows the 15-pin analog video connector on the rear of the engine. This connector conforms to the industry standard.

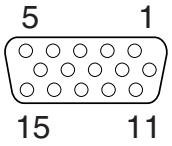

Related topics:

- "Input and output ports" on page 372
- "Starting the Configuration/Setup Utility" on page 106

Appendix E. Hardware messages and codes

This topic describes the hardware messages and codes associated with the SAN File System engine, which includes:

- "Beep symptoms"
- "Bus fault error messages" on page 380
- "Hard disk drive error messages" on page 381
- "Diagnostic error codes" on page 381
- "Fan error messages" on page 386
- "Host built-in self test (BIST) error messages" on page 387
- "Light path diagnostic panel LEDs" on page 387
- "POST error codes" on page 389
- "Power error messages" on page 397
- "Temperature error messages" on page 399
- "Temperature-related system shutdown messages" on page 400
- "Voltage-related system shutdown messages" on page 402

Beep symptoms

Beep symptoms are short tones or a series of short tones separated by pauses (intervals without sound). They occur when you power on an engine and can be used to determine whether or not the system is functioning properly.

Use this table to determine the components that need to be replaced based on the beep symptoms that you are experiencing. The following beep symptom examples show how to read the beep/symptom column in the table.

- 1-2-3 indicates:
  - One beep
  - A pause (or break)
  - Two beeps
  - A pause (or break)
  - Three beeps
- 4 indicates four continuous beeps.

One beep after successfully completing POST indicates the system if functioning properly.

Table 10. Beep symptoms

Beep/symptom	Part/action
1-1-2	The microprocessor register test failed. 1. Optional microprocessor 2. Microprocessor 3. System board
1-1-3	CMOS write/read test failed 1. Battery 2. System planar and shuttle assembly

Table 10. Beep symptoms (continued)

Beep/symptom	Part/action
1-1-4	BIOS EEPROM checksum failed 1. Recover BIOS 2. System board
1-2-1	Programmable Interval Timer failed. Replace the system board.
1-2-2	DMA initialization failed. Replace the system board.
1-2-3	DMA page register write/read failed. Replace the system board.
1-2-4	RAM refresh verification failed 1. DIMM 2. System board
1-3-1	First 64K RAM test failed. Replace the DIMM.
2-1-1	Secondary DMA register failed. Replace the system board.
2-1-2	Primary DMA register failed. Replace the system board.
2-1-3	Primary interrupt mask register failed. Replace the system board.
2-1-4	Secondary interrupt mask register failed. Replace the system board.
2-2-1	Interrupt vector loading failed. Replace the system board.
2-2-2	Keyboard controller failed 1. System board 2. Keyboard
2-2-3	CMOS power failure and checksum checks failed 1. Battery 2. System board
2-2-4	CMOS configuration information validation failed 1. Battery 2. System board
2-3-1	Screen initialization failed. Replace the system board.
2-3-2	Screen memory failed. Replace the system board.
2-3-3	Screen retrace failed. Replace the system board.
2-3-4	Search for video ROM failed. Replace the system board.
2-4-1	Video failed, but the monitor seems to be operable. Replace the system board.
3-1-1	Timer tick interrupt failed. Replace the system board.
3-1-2	Interval timer channel 2 failed. Replace the system board.
3-1-3	RAM test failed above address OFFFFH 1. DIMM 2. System board.
3-1-4	Time-of-day clock failed 1. Battery 2. System board
3-2-1	Serial port failed. Replace the system board.

Table 10. Beep symptoms (continued)

Beep/symptom	Part/action
3-2-2	Parallel port failed. Replace the system planar and system board.
3-2-3	Math coprocessor test failed 1. Microprocessor 2. System board.
3-2-3	Failure comparing CMOS memory size against actual 1. DIMM 2. Battery
3-3-1	Memory size mismatch occurred. 1. DIMM 2. Battery
3-3-2	Critical SMBUS error occurred. 1. Disconnect the server power cord from outlet, wait 30 seconds and retry. 2. System board. 3. DIMMs. 4. Hard disk drive backplane. 5. Power supply. 6. Power cage assembly. 7. I2C Cable.
3-3-3	No operational memory in system. 1. Install or reseat the memory modules. 2. Use the Configuration/Setup Utility to perform 3–boot reset. 3. DIMMs. 4. Memory board. 5. System board.
4-4-4	Optional system management adapter not installed in slot 1 or not functioning correctly. 1. Verify that the adapter is installed in slot 1. 2. Adapter. 3. System board.
Two short beeps	The configuration has changed (information only). 1. Run Diagnostics. 2. Run the Configuration/Setup Utility.
Three short beeps	1. DIMM 2. System board
One continuous beep	1. Microprocessor 2. Optional microprocessor 3. System board
Repeating short beeps	1. Keyboard 2. System board
One long and one short beep	1. Video adapter (if installed) 2. System planar and shuttle assembly

Table 10. Beep symptoms (continued)

Beep/symptom	Part/action
One long and two short beeps	1. Video adapter (if installed) 2. System board
One long and three short beeps	1. Monitor 2. Video adapter, if installed 3. System board
Two long and two short beeps	Replace the video adapter.

Related topics:

- Appendix F, "Parts listings", on page 405
- "Starting the Configuration/Setup Utility" on page 106
- "Running engine diagnostics" on page 104

Bus fault error messages

Table 11. Bus fault error messages

Message	Action
Failure reading I^2C device.	Check devices on bus 0. 1. Reseat the I^2C cable between the Remote Supervisor II Adapter (in PCI slot 5/J5 on the PCI riser card) and the remote supervisor adapter II connector on the PCI riser card. 2. Memory DIMMs. 3. System board.
Failure reading I^2C device.	Check devices on bus 1. 1. Reseat the I2C cable between the operator information panel and the system board (J22). 2. Operator information panel. 3. System board.
Failure reading I^2C device.	Check devices on bus 2. 1. Reseat the cable between the system board and the power supply (power cage assembly) (J10). 2. Power cage assembly. 3. Power supply. 4. System board.
Failure reading I^2C device.	Check devices on bus 3. 1. Reseat the cable between the hard disk drive backplane and the connector (J10) of the system board. 2. Hard disk drive backplane. 3. System board.
Failure reading I^2C device.	Check devices on bus 4. Replace the system board.

Hard disk drive error messages

Table 12. Hard disk drive error messages

Message	Action
Hard drive*x* removal detected	Hard drive*x* has been removed. Level: Critical Information only. Take action as appropriate.

Related topics:

- "Replacing the hard disk backplane" on page 120

Diagnostic error codes

Use the diagnostic error codes to determine the part to be replaced or action to be taken for correcting the problem indicated by the error code.

Note: In the following error codes, if *XXX* is *000*, *195*, or *197*, *do not* replace a FRU. The description for these error codes are:

000 The test passed.

195 The Esc key was pressed to stop the test.

197 Warning; a hardware failure might not have occurred.

See "Model 1RX" on page 405 to determine which components should be replaced by a field service technician.

Table 13. Diagnostics error codes

Error code/symptom	Part/action
001-XXX-000	Failed core tests. Replace the system board.
001-XXX-001	Failed core tests. Replace the system board.
001-250-000	Failed system board ECC. Replace the system board.
001-250-001	Failed system board ECC. Replace the system board.
005-XXX-000	Failed video test. Replace the system board.
011-XXX-000	Failed COM1 serial port test. Replace the system board.
011-XXX-001	Failed COM2 serial port test. Replace the system board.
014-XXX-000	Failed parallel port test. Replace the system board.
015-XXX-001	USB interface not found. The system board is damaged. Replace the system board.
015-XXX-015	Failed USB external loopback test. 1. Make sure the parallel port is not disabled. 2. Run the USB external loopback test again. 3. System board.
015-XXX-198	USB device connected during USB test. 1. Remove USB devices from USB1 and USB2. 2. Run the USB external loopback test again. 3. System board.
020-XXX-000	Failed PCI interface test. Replace the system board.

Table 13. Diagnostics error codes (continued)

Error code/symptom	Part/action
020-XXX-001	Failed hot-swap slot 1 PCI latch test. 1. PCI hot-swap latch assembly 2. System board
020-XXX-002	Failed hot-swap slot 2 PCI latch test. 1. PCI hot-swap latch assembly 2. System board
020-XXX-003	Failed hot-swap slot 3 PCI latch test. 1. PCI hot-swap latch assembly 2. System board
020-XXX-004	Failed hot-swap slot 4 PCI latch test. 1. PCI hot-swap latch assembly 2. System board
030-XXX-000	Failed internal SCSI interface test. Replace the system board.
035-XXX-099	No adapters were found. If an adapter is installed, re-check the connection.
035-XXX-S99	Failed RAID test on PCI slot *s*, where *s* is the number of failing PCI slot. Check system Error Log before replacing a FRU. 1. Check the system error log and resolve any problems. 2. Adapter 3. SCSI backplane 4. Cable
035-XXX-SNN	Failed RAID test on PCI slot *s*, where *s* is the number of failing PCI slot and *nn* is the SCSI ID of failing fixed disk. • Check the system error log and resolve any problems. • Hard disk drive with SCSI ID *nn* on RAID adapter in PCI slot *s*.
035-253-S99	RAID adapter initialization failure in slot *s*, where *s* is the number of the slot. 1. Obtain the basic and extended configuration status and see the *ServeRAID Hardware Maintenance Manual* for more information. 2. Cable. 3. SCSI backplane. 4. Adapter.
075-XXX-000	Failed power supply test. Replace the power supply.
089-XXX-001	Failed microprocessor test. 1. VRM 1 for microprocessor 1 2. Microprocessor 1
089-XXX-002	Failed optional microprocessor test. 1. VRM 2 for optional microprocessor 2 2. Optional microprocessor 2

Table 13. Diagnostics error codes (continued)

Error code/symptom	Part/action
166-198-000 System Management: Aborted	Unable to communicate with ASM. It may be busy. 1. Run the diagnostic test again. 2. Correct other error conditions and retry. These include other failed system management tests and items logged in the system error log of the Remote Supervisor Adapter II. 3. Disconnect all engine and option power cords from the engine, wait 30 seconds, reconnect, and retry. 4. Replace RSA II. 5. Replace System board.
166-201-001 System Management: Failed	One or more I^2C bus errors. 1. See SERVPROC and DIAGS entries in the event log. 2. Reseat the I^2C cable between the Remote Supervisor Adapter (in PCI slot 5/J5 on the PCI riser card) and system management connector on the PCI riser card. 3. Reseat memory DIMMs. 4. Replace memory DIMMs. 5. Replace System board.
166-201-002 System Management: Failed	One or more I^2C bus errors. 1. See SERVPROC and DIAGS entries in the event log 2. Reseat I^2C cable between the operator information panel and the system board (J22). 3. Replace diagnostics panel. 4. Replace system board.
166-201-003 System Management: Failed	One or more I^2C bus errors. 1. Reseat cables between the system board and the power supply or power cage assembly. 2. Replace power cage assembly. 3. Replace system board.
166-201-004 System Management: Failed	One or more I^2C bus errors. 1. See SERVPROC and DIAGS entries in the event log. 2. Replace hard disk drive backplane. 3. Replace system board.
166-201-005 System Management: Failed	One or more I^2C bus errors. 1. See SERVPROC and DIAGS entries in the event log. 2. Reseat memory DIMMs. 3. Reseat microprocessors. 4. Replace memory DIMMs. 5. Replace microprocessors. 6. Replace system board.

Table 13. Diagnostics error codes (continued)

Error code/symptom	Part/action
166-250-000 System Management: Failed	I^2C cable is disconnected. 1. Reconnect I^2C cable between Remote Supervisor Adapter II and system board. 2. Reseat the I^2C cable between the Remote Supervisor Adapter II (in PCI slot 5/J5 on the PCI riser card) and system management connector on the PCI riser card. 3. Replace I^2C cables. 4. Replace Advanced System Management (ASM) adapter. 5. Replace system board.
166-260-000 System Management: Failed	Restart ASM error. After restarting, ASM communication was lost. 1. Disconnect all engine and option power cords from the engine, wait 30 seconds, reconnect, and retry. 2. Reseat RSA II in PCI slot 5/J5 on the PCI riser card. 3. Replace RSA II.
166-342-000 System Management: Failed	ASM adapter BIST indicate failed tests. 1. Ensure the latest firmware levels for RSA II and BIOS. 2. Disconnect all engine and option power cords from engine, wait 30 seconds, reconnect, and retry. 3. Replace RSA II.
166-400-000 System Management: Failed	ISMP self test result failed tests: *x* where *x* is Flash, RAM, or ROM. 1. Reflash or update firmware for ISMP. 2. Replace system board.
180-XXX-000	Diagnostics LED failure. Run diagnostics panel LED test for the failing LED.
180-XXX-001	Failed front LED panel test. 1. Replace operator information panel. 2. Replace system board.
180-XXX-002	Failed diagnostics LED panel test. 1. Replace diagnostics panel. 2. Replace system board.
180-361-003	Failed fan LED test. Replace fan or fans.
180-XXX-003	Failed system board LED test. Replace system board.
180-XXX-005	Failed SCSI backplane LED test. 1. Replace SCSI backplane. 2. Replace SCSI backplane cable. 3. Replace system board.
201-XXX-0NN	Failed memory test, where *nn* is the slot location (slots 1–6). 1. DIMM Location slots 1-6 where *nn* is the DIMM location. **Note:** *nn* 1=DIMM 1; 2=DIMM 2; 3=DIMM 3; 4=DIMM 4. 2. Replace system board.

Table 13. Diagnostics error codes (continued)

Error code/symptom	Part/action
201-XXX-999	Multiple DIMM failure, see error text. 1. See error text for failing DIMMs. 2. Replace system board.
202-XXX-001	Failed system cache test. 1. Replace VRM 1 2. Replace microprocessor 1
202-XXX-002	Failed system cache test. 1. Replace VRM 2. 2. Replace microprocessor 2.
206-XXX-000	Failed diskette drive test. 1. Run the test again using a different diskette. 2. Replace cable. 3. Replace diskette drive. 4. Replace system board.
215-XXX-000	Failed IDE CD-ROM drive test. 1. Run the test again with a different CD-ROM. 2. Replace CD-ROM drive cables. 3. Replace CD-ROM drive. 4. Replace system board.
217-198-XXX	Could not establish drive parameters. 1. Check cable and termination. 2. Replace SCSI backplane. 3. Replace hard disk.
217-XXX-000	Failed BIOS hard disk test. **Note:** If RAID is configured, the hard disk number refers to the RAID logical array. Replace hard disk 1.
217-XXX-001	Failed BIOS hard disk test. **Note:** If RAID is configured, the hard disk number refers to the RAID logical array. Replace hard disk 2.
217-XXX-002	Failed BIOS hard disk test. **Note:** If RAID is configured, the hard disk number refers to the RAID logical array. Replace hard disk 3.
217-XXX-003	Failed BIOS hard disk test. **Note:** If RAID is configured, the hard disk number refers to the RAID logical array. Replace hard disk 4.
217-XXX-004	Failed BIOS hard disk test. **Note:** If RAID is configured, the hard disk number refers to the RAID logical array. Replace hard disk 5.

Table 13. Diagnostics error codes (continued)

Error code/symptom	Part/action
217-XXX-005	Failed BIOS hard disk test. **Note:** If RAID is configured, the hard disk number refers to the RAID logical array. Replace hard disk 6.
264-XXX-0NN	Failed tape drive test 1. Tape cartridge, if you executed the Read/Write Tape Drive test (failure code of XXX = 256). 2. Replace SCSI or power cable connected to tape drive with SCSI ID *nn*. 3. Replace Tape drive with SCSI ID *nn* (refer to the Help and Service Information appendix of the tape drive's User Guide). 4. Replace system board or SCSI controller (run SCSI controller diagnostic to determine if the SCSI bus is functioning properly.).
264-XXX-999	Errors on multiple tape drives. See error messages/text in the PC Doctor error log for detailed information on each individual tape drive error.
301-XXX-000	Failed keyboard test. Replace keyboard.
405-XXX-000	Failed Ethernet test on controller on the system board. 1. Verify that Ethernet is not disabled in BIOS. 2. Replace system board.
405-XXX-00N	Failed Ethernet test on adapter in PCI slot *n*. 1. Adapter in PCI slot *n*. 2. Replace system board.
415-XXX-000	Failed modem test. 1. Replace cable. **Note:** Ensure modem is present and attached to engine. 2. Replace modem. 3. Replace system board.

Related topics:

• Appendix F, "Parts listings", on page 405

Fan error messages

Table 14. Fan error messages

Message	Action
Fax *x* failure	Fan *x* failed. Level: Critical 1. Check all connections to fan *x*. 2. Replace fan *x*.

Table 14. Fan error messages (continued)

Message	Action
Fan *x* fault	Fan *x* operating beyond recommended RPM range. Level: Critical 1. Check all connections to fan *x*. 2. Replace fan *x*.
Fax *x* outside recommended speed action	Fan *x* had an under-temperature condition. Level: Warning Replace fan *x*

Related topics:

- "System reliability considerations" on page 425
- "Hardware specifications" on page 6
- "Replacing a hot-swap fan" on page 123

Host built-in self test (BIST) error messages

Table 15. Host build-in self test error messages

Message	Action
Host fail	The built-in self test for the host failed. Level: Informational 1. Reseat the microprocessor. 2. Reseat the VRM. 3. Replace the microprocessor CPU.

Related topics:

- "Replacing a microprocessor" on page 131

Light path diagnostic panel LEDs

If the system-error LED on the front of the server is lit, check the Light Path Diagnostics panel to see whether any LEDs on that panel are lit.

Table 16. Light Path Diagnostics panel LEDs

LED lit	Action
No LEDs are lit	Either the system error log is at least 75% full or a Predictive Failure Analysis® (PFA) alert was logged. View the system error log and correct any problems. **Note:** Disconnect the engine from all power sources for at least 20 seconds to turn off the system-error LED.

Table 16. Light Path Diagnostics panel LEDs (continued)

LED lit	Action
CPU	One of the microprocessors has failed or a microprocessor is installed in the wrong socket. 1. Check the microprocessor-error LEDs on the system board. If a microprocessor-error LED is lit for a microprocessor socket that is empty, the microprocessors are not installed in the correct order. 2. Power off the engine, reseat the microprocessor indicated by the lit microprocessor-error LED, and restart the engine. 3. If the problem persists, replace the microprocessor.
VRM	One of the microprocessor Voltage Regulator Modules (VRMs) has failed. 1. Check the microprocessor VRM connectors on the system board to locate the error LED next to the failing component. Then, replace the failing component. 2. Power OFF the engine, reseat the microprocessor VRM indicated by the lit VRM-error LED, and restart the engine. 3. If the problem persists, replace the microprocessor VRM.
MEMORY	A memory error occurred. 1. Check the DIMM error LEDs on the system board. 2. Replace the DIMM indicated by the lit DIMM-error LED.
PCI BUS A **PCI BUS B** **PCI BUS C**	An error occurred on PCI bus A, B, or C. An adapter in PCI slot 1, 2, 3, 4, or 5 or the system board caused the error. 1. Check the adapter slots to locate the error LED next to the failing bus. 2. View the system-error log and correct any problems. 3. If the problem persists, try to determine the failing adapter by removing one adapter at a time from PCI bus A (PCI-X slots 1 or 2), PCI bus B (PCI-X slots 3 or 4), or PCI bus C (PCI slot 5). Restart the engine after removing each adapter.
DASD	A hot-swap hard disk drive has failed on internal SCSI channel A. 1. View the system error log and correct any problems. 2. If the system error log indicates a temperature problem and the fans are working correctly, verify that the ambient temperature is within normal limits. 3. If the hard disk status LED on one of the hot-swap disk drives is lit continuously, the drive has failed.
NMI	A nonmaskable interrupt occurred. If the PCI BUS LED is also lit, follow the actions dictated for the PCI BUS LED. If the PCI BUS LED is not lit, restart the engine.
ISMP	The integrated system management processor (ISMP) detected an internal error. 1. Updated the ISMP firmware to the latest level code. Then, unplug the AC power, wait at least 30 seconds, and restart the engine. 2. If the problem persists, replace the system board.
POWER SUPPLY 1	The power supply in bay 1 has failed. Replace the power supply.
POWER SUPPLY 2	The power supply in bay 2 has failed. Replace the power supply.

Table 16. Light Path Diagnostics panel LEDs (continued)

LED lit	Action
NON-REDUNDANT	The engine is operating in a nonredundant power mode for one of these reasons: • One power supply has failed or is not connected to a functioning AC electrical outlet. • The system has exceeded the power capabilities of one of the redundant power supplies. To resolve the problem: 1. Check the power supply LEDs to determine the failing power supply. After identifying the failing power supply: a. Verify the power supply connections. b. If a power supply has failed, replace the failing power supply. 2. If neither power supply LED is lit, remove any optional devices from the engine to restore redundancy.
FAN	One of the fans has failed or is operating too slowly. **Note:** A failing fan can also cause the TEMP and DASD LEDs to be on. 1. Remove the cover to determine which fan is failing. The LED on the failing fan will be lit. 2. Replace the fan.
TEMP	The system temperature has exceeded the maximum rating. 1. If the FAN LED is lit, perform the actions for that LED. 2. Make sure the room temperature is not too high.
CPU/VRM MISMATCH	The service processor detected a microprocessor or VRM mismatch. 1. If a microprocessor is missing from socket 1, install a microprocessor in socket 1. 2. Switch the microprocessor in socket 1 with the microprocessor in socket 2. 3. Verify that both microprocessors are the same type with the same core frequency and same L2 size. 4. Verify that both VRMs are the same type and that both VRMs are supported in the engine.

Related topics:

- "Replacing a microprocessor" on page 131
- "Hardware specifications" on page 6
- "Replacing the system planar and shuttle" on page 142

POST error codes

When you power-on an engine, it performs a series of tests, called the power-on self test (POST), that check the operation of each of the components as well as some of the options that are installed in the engine. If POST detects a problem, an error message appears on the screen (and you will hear a series of one or more beeps). Use this table to determine how to resolve problems using the POST error codes. You can also view POST error messages in the POST error log.

One beep after successfully completing POST indicates the system if functioning properly.

Table 17. POST error codes

Error code/symptom	Part/action
062	Three consecutive startup failures using the default configuration. 1. Run the Configuration/Setup Utility. 2. Replace battery. 3. Replace system board. 4. Replace microprocessor.
101, 102	System and processor error. Replace system board.
106	System and processor error. Replace system board.
111	Channel check error 1. Replace memory DIMM. 2. Replace system board.
114	Adapter read-only memory error 1. Replace failing adapter. 2. Run diagnostics.
129	Internal cache error 1. Replace microprocessor. 2. Replace optional microprocessor.
151	Real time clock error 1. Run diagnostics. 2. Replace battery. 3. Replace system board.
161	Real time clock battery error 1. Run the Configuration/Setup Utility. 2. Replace battery. 3. Replace system board.
162	Device configuration error **Note:** Be sure to load the default settings and any additional desired settings; then, save the configuration. 1. Run the Configuration/Setup Utility. 2. Replace battery. 3. Replace failing device. 4. Replace system board.
163	Real-time clock error 1. Run the Configuration/Setup Utility. 2. Replace battery. 3. Replace system board.
164	Memory configuration changed. 1. Run the Configuration/Setup Utility. 2. Replace DIMM. 3. Replace system board.
175	Hardware error. Replace system board.

Table 17. POST error codes (continued)

Error code/symptom	Part/action
176	Computer cover or cable cover was removed without a key being used. 1. Run the Configuration/Setup Utility. 2. Replace system board.
177, 178	Security hardware error 1. Run the Configuration/Setup Utility. 2. Replace system board.
184	Power-on password damaged. 1. Run the Configuration/Setup Utility. 2. Replace system board.
185	Drive startup sequence information corrupted. 1. Run the Configuration/Setup Utility. 2. Replace system board.
186	Security hardware control logic failed. 1. Run the Configuration/Setup Utility. 2. Replace system board.
187	VPD serial number not set. 1. Run the Configuration/Setup Utility to set the serial number. 2. Replace system board.
188	Bad EEPROM CRC number 2 1. Run the Configuration/Setup Utility. 2. Replace system board.
189	An attempt was made to access the engine with passwords that were not valid. Run the Configuration/Setup Utility and type the administrator password.
201	Memory test error 1. If the engine does not have the latest level of BIOS installed, update the BIOS to the latest level and run the diagnostic program again. 2. Replace DIMM. 3. Replace system board.
229	Cache error 1. Replace microprocessor. 2. Replace optional microprocessor.
262	DRAM parity configuration error. 1. Run the Configuration/Setup Utility. 2. Replace battery. 3. Replace system board.
289	DIMM disabled by POST or user. 1. Run the Configuration/Setup Utility, if the DIMM was disabled by the user. 2. Replace disabled DIMM, if not disabled by user.

Table 17. POST error codes (continued)

Error code/symptom	Part/action
301	Keyboard or keyboard controller error 1. Replace keyboard. 2. Replace system board.
303	Keyboard controller error. Replace system board.
602	Diskette boot record is not valid. 1. Replace diskette. 2. Replace diskette drive. 3. Replace cable. 4. Replace system board.
604	Diskette drive error 1. Run the Configuration/Setup Utility and diagnostics. 2. Replace diskette drive. 3. Replace drive cable. 4. Replace system board.
605	Unlock failure 1. Replace diskette drive. 2. Replace drive cable. 3. Replace system board.
662	Diskette drive configuration error. 1. Run the Configuration/Setup Utility and diagnostics. 2. Replace diskette drive. 3. Replace drive cable. 4. Replace system board.
762	Coprocessor configuration error. 1. Run the Configuration/Setup Utility. 2. Replace battery. 3. Replace microprocessor.
962	Parallel port error 1. Disconnect the external cable on the parallel port. 2. Run the Configuration/Setup Utility. 3. Replace system board.
11XX	System board serial port 1 or 2 error 1. Disconnect the external cable on the serial port. 2. Run the Configuration/Setup Utility. 3. Replace system board.
1301	I^2C cable to front panel not found 1. Replace cable. 2. Replace front panel. 3. Replace power switch assembly. 4. Replace system board.

Table 17. POST error codes (continued)

Error code/symptom	Part/action
1302	I^2C cable from system board to power on and reset switches not found. 1. Replace cable. 2. Replace power switch assembly. 3. Replace system board.
1303	I^2C cable from system board to power backplane not found. 1. Replace cable. 2. Replace power cage assembly. 3. Replace system board.
1304	I^2C cable to diagnostic LED board not found. 1. Replace power switch assembly. 2. Replace system board.
1600	The system management processor is not functioning. 1. Ensure that a jumper is not installed on J34. 2. Remove the AC power to the engine and wait 20 seconds. Then reconnect the AC power. Wait 30 seconds and then power ON the engine. 3. Replace system board.
1601	The system is able to communicate to the system management processor, but the system management processor failed to respond at the start of POST. 1. Remove the AC power to the engine and wait 20 seconds. Then reconnect the AC power. Wait 30 seconds; then, power ON the engine. 2. Flash update the system management processor. 3. Replace RSA II. 4. Replace system board.
1602	Cable for optional service processor adapter not installed. Disconnect all engine and option power cords from engine. Wait 30 seconds, reconnect, and retry.
1762 (Hard disk configuration error)	1. Replace hard disk drive. 2. Replace hard disk cables. 3. Run the Configuration/Setup Utility. 4. Replace hard disk adapter. 5. Replace SCSI backplane. 6. Replace system board.
178X	Fixed disk error 1. Replace hard disk cables. 2. Run diagnostics. 3. Replace hard disk adapter. 4. Replace hard disk drive. 5. Replace system board.

Table 17. POST error codes (continued)

Error code/symptom	Part/action
1800	No more hardware interrupt available for PCI adapter. 1. Run the Configuration/Setup Utility. 2. Replace failing adapter. 3. Replace system board.
1962	Drive does not contain a valid boot sector. 1. Verify that a startable operating system is installed. 2. Run diagnostics. 3. Replace hard disk drive. 4. Replace SCSI backplane. 5. Replace cable. 6. Replace system board.
2400	Video controller test failure. 1. Replace video adapter. 2. Replace system board.
2462	Video memory configuration error. 1. Replace video adapter. 2. Replace system board.
5962	IDE CD-ROM drive configuration error. 1. Run the Configuration/Setup Utility. 2. Replace CD-ROM drive. 3. Replace CD-ROM power cable. 4. Replace IDE cable. 5. Replace system board. 6. Replace battery.
8603	Pointing-device error 1. Replace pointing device. 2. Replace system board.
0001200	Machine check architecture error. 1. Replace microprocessor 1. 2. Replace optional microprocessor 2.
00012000	Microprocessor machine check. 1. Replace microprocessor. 2. Replace system board.
00019501	Microprocessor 1 is not functioning. 1. Check VRM and microprocessor LEDs. 2. Replace VRM 1. 3. Replace microprocessor 1. 4. Replace system board.
00019502	Microprocessor 2 is not functioning . 1. Check VRM and microprocessor LEDs. 2. Replace VRM 2. 3. Replace microprocessor 2.

Table 17. POST error codes (continued)

Error code/symptom	Part/action
00019701	Microprocessor 1 failed. 1. Replace microprocessor 1. 2. Replace system board.
00019702	Microprocessor 2 failed. 1. Replace microprocessor 2. 2. Replace system board.
00180100	A PCI adapter has requested memory resources that are not available. 1. Reorder the adapters in the PCI slots. It is important that your startup device is positioned early in the startup-device order so that it is run by POST. 2. Ensure that the PCI adapter and all other adapters are set correctly in the Configuration/Setup Utility. If the memory resource settings are not correct, change the settings. 3. If all memory resources are being used, you might need to remove an adapter to make memory available to the PCI adapter. Disabling the adapter BIOS on the adapter might correct the error. (See the documentation provided with the adapter.)
00180200	No more I/O space available for PCI adapter. 1. Run the Configuration/Setup Utility. 2. Replace failing adapter. 3. Replace system board.
00180300	No more memory (above 1 MB for PCI adapter). 1. Run the Configuration/Setup Utility. 2. Replace failing adapter. 3. Replace system board.
00180400	No more memory (below 1 MB for PCI adapter). 1. Run the Configuration/Setup Utility. 2. Move the failing adapter to slot 1. 3. Replace failing adapter. 4. Replace system board.
00180500	PCI option ROM checksum error. 1. Remove failing PCI card. 2. Replace system board.
00180600	PCI-to-PCI bridge error. 1. Run the Configuration/Setup Utility. 2. Move the failing adapter to slot 1. 3. Replace failing adapter. 4. Replace system board.
00180700, 00180800	General PCI error. 1. Replace system board. 2. Replace PCI card.

Table 17. POST error codes (continued)

Error code/symptom	Part/action
00181000	PCI error. • Replace adapter. • Replace system board.
01295085	ECC checking hardware test error. 1. Replace system board. 2. Replace microprocessor.
01298001	No update data for microprocessor 1 1. Ensure all processors have the same cache size. 2. Replace microprocessor 1.
01298002	No update data for microprocessor 2 1. Ensure all processors have the same cache sizes, dock speeds and clock frequencies. 2. Replace microprocessor 2.
01298101	Bad update data for microprocessor 1. 1. Ensure all processors have the same cache sizes, dock speeds and clock frequencies. 2. Replace microprocessor 1.
01298102	Bad update data for microprocessor 2. 1. Ensure all processors have the same cache sizes, dock speeds and clock frequencies. 2. Replace microprocessor 2.
I9990301	Hard disk sector error 1. Replace hard disk drive. 2. Replace SCSI backplane. 3. Replace cable. 4. Replace system board.
I9990305	Hard disk sector error, no operating system installed. Install operating system to hard disk.
I9990650	AC power has been restored. 1. Check cable. 2. Check for interruption of power. 3. Replace power cable.

Related topics:

- Appendix F, “Parts listings”, on page 405
- Chapter 11, “Adding and replacing engine components”, on page 109

Power error messages

Table 18. Power error messages

Message	Action
Power supply*x* **current share fault**	There is excessive current demand on power supply *x*. Level: Critical Replace power supply *x*.
Power supply *x* **DC good fault**	A power good signal is not detected for power supply*x*. Level: Critical Replace power supply *x*.
Power supply*x* **temperature fault**	Power supply *x* had an over-temperature condition. Level: Critical Replace power supply *x*.
Power supply*x* **removed**	Power supply *x* has been removed. Level: Informational No action is required.
Power supply*x* **fan fault**	There is a fan fault in power supply *x*. Level: Critical Replace power supply *x*.
Power supply*x* **12 V fault**	An overcurrent condition has been detected. Level: Critical 1. Power off the engine. 2. Disconnect all AC power cords. 3. Check for loose cables in power subsystem. Also check for short circuits. For example, check to see if there is a loose screw causing a short circuit on a circuit board. 4. Remove adapters. Disconnect cables and power connectors to all internal and external devices until the engine is at the minimum configuration required to power ON the engine. **Note:** Minimum operating requirements for the engine are: a. One power supply b. Power cage assembly c. System board d. One microprocessor and VRM e. Memory module (with a minimum of two 128-MB DIMMs) 5. Reconnect all AC power cords and power ON the engine. If the engine starts successfully, replace adapters and devices one at a time until the problem is isolated. If the engine does not start with the minimum configuration, replace each of the components in the minimum configuration one at a time until the problem is isolated.

Table 18. Power error messages (continued)

Message	Action
Power supply*x* 3.3 V fault	3.3 V power supply *x* had an error. Level: Critical 1. Power off the engine. 2. Disconnect all AC power cords. 3. Check for loose cables in power subsystem. Also check for short circuits. For example, check to see if there is a loose screw causing a short circuit on a circuit board. 4. Remove adapters. Disconnect cables and power connectors to all internal and external devices until the engine is at the minimum configuration required to power ON the engine. **Note:** Minimum operating requirements for the engine are: a. One power supply b. Power cage assembly c. System board d. One microprocessor and VRM e. Memory module (with a minimum of two 128-MB DIMMs) 5. Reconnect all AC power cords and power ON the engine. If the engine starts successfully, replace adapters and devices one at a time until the problem is isolated. If the engine does not start with the minimum configuration, replace each of the components in the minimum configuration one at a time until the problem is isolated.
Power supply*x* 5 V fault	5 V power supply *x* had an error. Level: Critical 1. Power off the engine. 2. Disconnect all AC power cords. 3. Check for loose cables in power subsystem. Also check for short circuits. For example, check to see if there is a loose screw causing a short circuit on a circuit board. 4. Remove adapters. Disconnect cables and power connectors to all internal and external devices until the engine is at the minimum configuration required to power ON the engine. **Note:** Minimum operating requirements for the engine are: a. One power supply b. Power cage assembly c. System board d. One microprocessor and VRM e. Memory module (with a minimum of two 128-MB DIMMs) 5. Reconnect all AC power cords and power ON the engine. If the engine starts successfully, replace adapters and devices one at a time until the problem is isolated. If the engine does not start with the minimum configuration, replace each of the components in the minimum configuration one at a time until the problem is isolated.
System running non-redundant power	System does not have redundant power. Level: Noncritical System can continue to operate without redundancy protection. However, to add redundancy protection: 1. Add another power supply. 2. Remove options from the system.

Table 18. Power error messages (continued)

Message	Action
System under recommended voltage for *x*	The indicated voltage supply is under nominal value, where *x* is +12, −12, or +5. Level: Critical 1. Check all connections to the power subsystem. 2. Replace the power supply. 3. Replace the power cage assembly.

Related topics:

- "Troubleshooting power problems" on page 97
- "Replacing a hot-swap power supply" on page 125
- "System reliability considerations" on page 425

Temperature error messages

Table 19. Temperature error messages

Message	Action
DASD over temperature	Direct access storage device bay *x* was over temperature or sensor for DASD1 reported temperature over recommended range. Level: Critical Make sure that the engine is being properly cooled. ("System reliability considerations" on page 425)
DASD over recommended temperature	Direct access storage device bay *x* had an over-temperature condition. Level: Warning Make sure that the engine is being properly cooled. ("System reliability considerations" on page 425)
DASD under recommended temperature	Direct access storage device bay *x* had an under-temperature condition. Level: Warning Make sure that ambient temperature is within normal operating specifications. ("Hardware specifications" on page 6)
Power supply *x* **temperature fault**	Power supply *x* had an over-temperature condition. Level: Critical 1. Ensure that the engine is being properly cooled. ("System reliability considerations" on page 425) 2. Replace power supply *x*.
System board is over recommended temperature	The system board had an over-temperature condition. Level: Warning 1. Ensure that the engine is being properly cooled. ("System reliability considerations" on page 425) 2. Replace the system board.

Table 19. Temperature error messages (continued)

Message	Action
System board is under recommended temperature	The system board had an under-temperature condition. Level: Warning Make sure that ambient temperature is within normal operating specifications. (“Hardware specifications” on page 6)
System over temperature for CPU *x*	CPU *x* reported an over-temperature condition. Level: Warning Ensure that the engine is being properly cooled. (“System reliability considerations” on page 425)
System under recommended CPU *x***temperature**	The system is reporting an under-temperature condition for CPU *x* Level: Warning Make sure that ambient temperature is within normal operating specifications. (“Hardware specifications” on page 6)

Related topics:

- “System reliability considerations” on page 425
- “Hardware specifications” on page 6
- “Replacing a hot-swap power supply” on page 125
- “Replacing the system planar and shuttle” on page 142

Temperature-related system shutdown messages

Table 20. Temperature-related system shutdown error messages

Message	Action
System shutoff due to board over temperature	The system board is over temperature. Level: Critical 1. Ensure that the system is being properly cooled. See “System reliability considerations” on page 425. 2. Replace the system board.
System shutoff due to CPU *x* **over temperature**	CPU *x* is over temperature. Level: Critical 1. Ensure that the system is being properly cooled. See “System reliability considerations” on page 425. 2. Replace CPU *x*.
System shutoff due to CPU *x* **under temperature**	CPU *x* is under temperature. Level: Critical Ambient temperature must be within normal operating specifications. See “Hardware specifications” on page 6.

Table 20. Temperature-related system shutdown error messages (continued)

Message	Action
System shutoff due to DASD temperature (sensor*x*)	DASD area reported temperature outside recommended operating range. Level: Critical Ensure that the system is being properly cooled. See "System reliability considerations" on page 425.
System shutoff due to high ambient temperature.	High ambient temperature. Level: Critical Ambient temperature must be within normal operating specifications. See "Hardware specifications" on page 6.
System shutoff due to system board under temperature.	System board is under temperature. Level: Critical Ambient temperature must be within normal operating specifications. See "Hardware specifications" on page 6.

Related topics:

- "Hardware specifications" on page 6
- "System reliability considerations" on page 425
- "Replacing the system planar and shuttle" on page 142
- "Replacing a microprocessor" on page 131

Voltage-related system shutdown messages

Table 21. Voltage-related system shutdown error messages

Message	Action
System shutoff due to *x* current over max	The system is drawing too much current on voltage*x* bus. Level: Critical 1. Power off the engine. 2. Disconnect all AC power cords. 3. Check for loose cables in power subsystem. Also check for short circuits. For example, check to see if there is a loose screw causing a short circuit on a circuit board. 4. Remove adapters. Disconnect cables and power connectors to all internal and external devices until the engine is at the minimum configuration required to power ON the engine. **Note:** Minimum operating requirements for the engine are: a. One power supply b. Power cage assembly c. System board d. One microprocessor and VRM e. Memory module (with a minimum of two 128-MB DIMMs) 5. Reconnect all AC power cords and power ON the engine. If the engine starts successfully, replace adapters and devices one at a time until the problem is isolated. If the engine does not start with the minimum configuration, replace each of the components in the minimum configuration one at a time until the problem is isolated.
System shutoff due to*x* V over voltage	The system shut off due to *x* supply over voltage. Level: Critical 1. Check all power-supply connectors. 2. Replace the power supply. 3. Replace the power cage assembly.
System shutoff due to*x* V under voltage	The system shut off due to *x* supply under voltage. Level: Critical 1. Check all power-supply connectors. 2. Replace the power supply. 3. Replace the power cage assembly.
System shutoff due to VRM*x* over voltage.	The system shut off due to VRM *x* over voltage. Level:Replace VRM *x*.

Table 21. Voltage-related system shutdown error messages (continued)

Message	Action
System shutoff due to excessive (<240 VA) loading.	1. Power off the engine. 2. Disconnect all AC power cords. 3. Check for loose cables in power subsystem. Also check for short circuits. For example, check to see if there is a loose screw causing a short circuit on a circuit board. 4. Remove adapters. Disconnect cables and power connectors to all internal and external devices until the engine is at the minimum configuration required to power ON the engine. **Note:** Minimum operating requirements for the engine are: a. One power supply b. Power cage assembly c. System board d. One microprocessor and VRM e. Memory module (with a minimum of two 128-MB DIMMs) 5. Reconnect all AC power cords and power ON the engine. If the engine starts successfully, replace adapters and devices one at a time until the problem is isolated. If the engine does not start with the minimum configuration, replace each of the components in the minimum configuration one at a time until the problem is isolated. 6. Cycle AC power on or off.

Related topics:

- "Replacing the power-supply cage" on page 137
- "Replacing a hot-swap power supply" on page 125

Appendix F. Parts listings

Parts listings are available for the following components:

- "Model 1RX"
- "Power cords" on page 407

Model 1RX

The following reference drawing contains a corresponding index for all field-replaceable and customer-replaceable units (parts) for the SAN File System Model 1RX. The index provides the part number, the quantity required (units), and a description of the part.

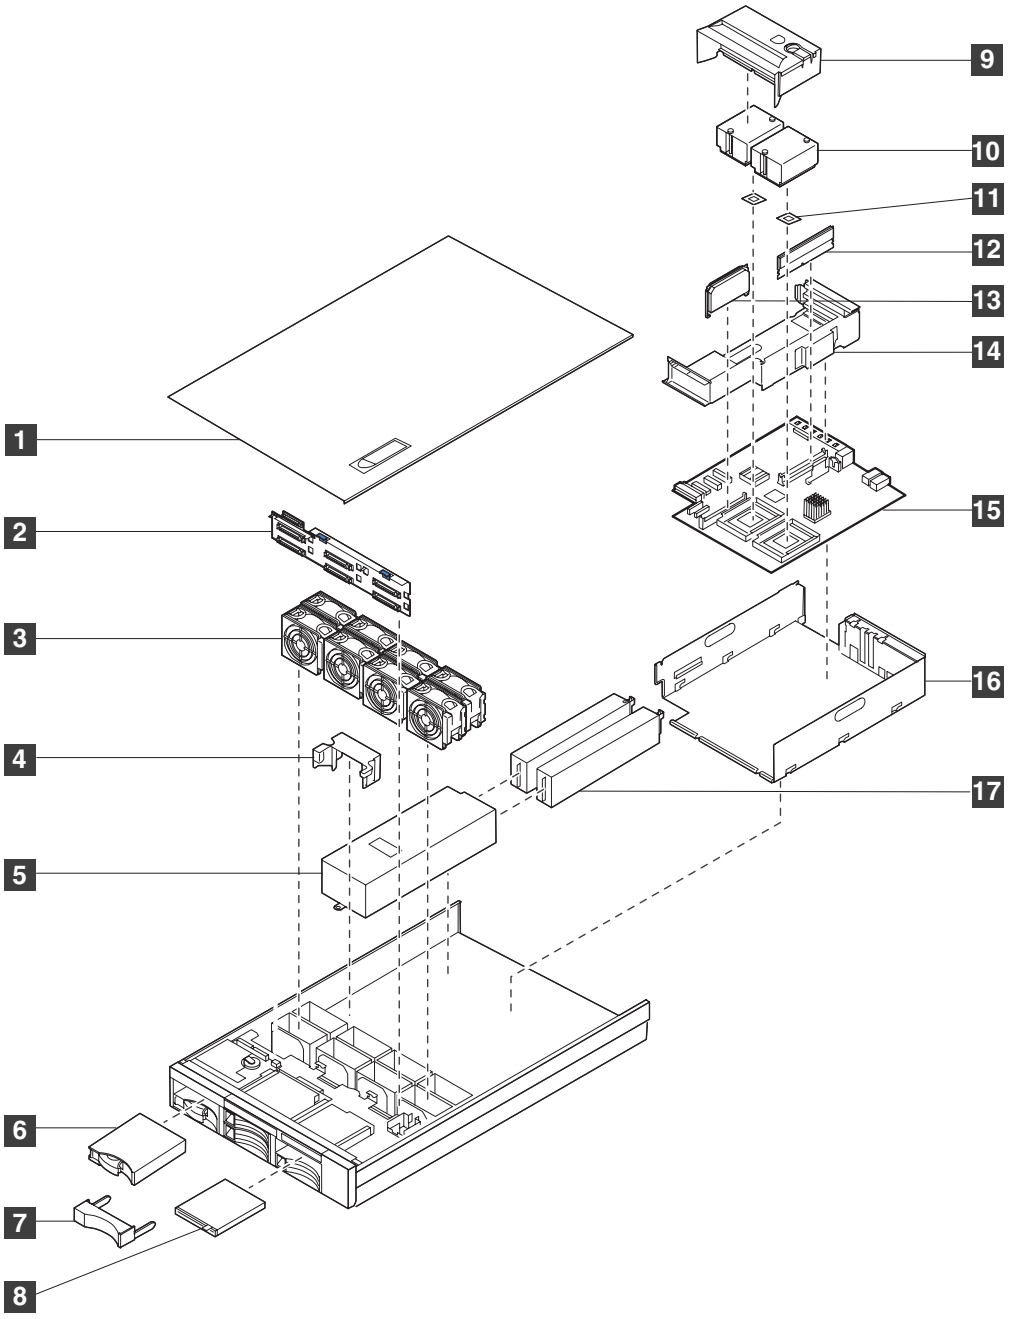

Table 22. Model 1RX parts listing

Index	Part number	Units	Description
1–	18P6148	2– 8	SAN File System Model 1RX
–1	18P6167	1	Top cover assembly
–2	59P5856	1	DASD backplane with bracket assembly
–3	01R0587	8	Fan assembly, 80-mm hot-swap
–4 and 9	01R0614	1	Thermal baffle kit (power-supply air baffle and microprocessor air baffle)
–5	49P2169	1	• Power cage and backplane assembly
–6	06P5759	2	• 36-GB 10K RPM hot-swap hard disk drive
–7	00N7259	4	• Hard disk drive filler blank

Table 22. Model 1RX parts listing (continued)

Index	Part number	Units	Description
–8			• CD-ROM drives
	06P5263		• • CD-ROM drive, 24X, primary
	06P5151		• • CD-ROM drive, 24X, alternate
	19K1523		• • CD-ROM drive, 24X, alternate
	33P3231		• • CD-ROM drive, 24X, alternate
–9 and 4	01R0614	1	• Thermal baffle kit (power-supply air baffle and microprocessor air baffle)
–10 and 11	71P8340	2	• Microprocessor and heat sink assembly
–12	09N4308	4	• Memory, 1-GB DDR SDRAM DIMM
–13	49P2010	2	• VRM, microprocessor
–14	73P6591	1	• PCI riser card assembly
–15	23K4456	1	• System board shuttle assembly
–16	49P2167	2	• Power supply, ac 514 watt, hot-swap
	48P9028		• CD interposer card assembly
	33F8354		• System battery
	01R0591		• Cable-management-arm assembly
			• Cables
	32P1869		• • CD signal to system board
	32P1874		• • CD signal interposer to system board
	32P1873		• • Fan 2 drop
	32P1870		• • Operator information card to system board
	06P5864		• • RSA 20–pin/TWP ribbon
	09N9581		• • RSA RS485
	32P1871		• • SCSI signal
	32P1872		• • USB
	36L8645		• Diskette drive, slim high
	01R0593		• EIA bracket assembly
	31P6027		• Thermal grease kit for heat sink
	6952300	1	• Power cord to wall outlet, 2.8 meters
	36L8886	1	• Power cord to power distribution unit (PDU), 2.8 meters
	24P8175	1	• 2-port 1/2-Gb/s Fibre Channel adapter
	22P7809	2	• 1-port 1-Gb/s Fiber Ethernet adapter
	73P9765	1	• Remote Service Adapter II (RSA II)
	59P2951	1	• RSA terminator
	01R0592	1	• Slide assembly
	48P9029	1	• Operator information card
	18P6152	1	• System service label
	73P6127	1	• Rack hardware kit

Related topics:

- Appendix F, “Parts listings”, on page 405
- “Power cords”

Power cords

IBM® power cords used in the United States and Canada are listed by Underwriter’s Laboratories (UL) and certified by the Canadian Standards Association (CSA).

For units intended to be operated at 115 volts: Use a UL-listed and CSA-certified cord set consisting of a minimum 18 AWG, Type SVT or SJT, three-conductor cord, a maximum of 15 feet in length and a parallel blade, grounding-type attachment plug rated 15 amperes, 125 volts.

For units intended to be operated at 230 volts (U.S. use): Use a UL-listed and CSA-certified cord set consisting of a minimum 18 AWG, Type SVT or SJT, three-conductor cord, a maximum of 15 feet in length and a tandem blade, grounding-type attachment plug rated 15 amperes, 250 volts.

For units intended to be operated at 230 volts (outside the U.S.): Use a cord set with a grounding-type attachment plug. The cord set should have the appropriate safety approvals for the country in which the equipment will be installed.

IBM power cords for a specific country or region are usually available only in that country or region.

Table 23. Power cords

Part number	Description
13F9940	Argentina, Australia, China (PRC), Colombia, Papua New Guinea, Paraguay, Uruguay, Western Samoa, New Zealand
13F9979	Afghanistan, Algeria, Andorra, Angola, Aruba, Austria, Belgium, Benin, Brazil, Bulgaria, Burundi, Cameroon, Central African Rep., Chad, Congo Brazzaville, Curacao, Czech Republic, Egypt, Finland, France, French Guiana, Germany, Greece, Guinea, Hungary, Iceland, Indonesia, Iran, Ivory Coast, Jordan, Kenya, Korea (South), Lebanon, Luxembourg, Macau, Malagasy, Mali, Martinique, Mauritania, Mauritius, Monaco, Morocco, Mozambique, Netherlands, Netherlands Antilles, New Caledonia, Niger, Norway, Poland, Portugal, Romania, Saudi Arabia, Senegal, Spain, Sudan, Sweden, Syria, Togo, Tunisia, Upper Volta, Russia, Vietnam, former Yugoslavia, Zaire, Zimbabwe
13F9997	Denmark
14F0015	Bangladesh, Burma, Pakistan, South Africa, Sri Lanka
14F0033	Antigua, Bahrain, Bermuda, Brunei, Channel Islands, China (Hong Kong S.A.R.), Cyprus, Fiji, Ghana, Guyana, India, Iraq, Ireland, Jordan, Kenya, Kuwait, Malawi, Malaysia, Malta, Nepal, Nigeria, Oman, Polynesia, Qatar, Sierra Leone, Singapore, Tanzania, Uganda, United Arab Emirates, United Kingdom, Yemen, Zambia
14F0051	Liechtenstein, Switzerland
14F0069	Chile, Ethiopia, Italy, Libya, Somalia
14F0087	Israel
1838574	Aruba, Bahamas, Barbados, Bermuda, Bolivia, Brazil, Canada, Cayman Islands, Costa Rica, Curacao, Dominican Republic, El Salvador, Ecuador, Guatemala, Guyana, Haiti, Honduras, Jamaica, Japan, Liberia, Netherlands Antilles, Nicaragua, Panama, Peru, Philippines, Suriname, Taiwan, Thailand, Trinidad Tobago, United States of America, Venezuela
6952300	Bahamas, Barbados, Bermuda, Bolivia, Brazil, Canada, Cayman Islands, Colombia, Costa Rica, Curacao, Dominican Republic, Ecuador, El Salvador, Guatemala, Guyana, Haiti, Honduras, Jamaica, Japan, Korea (South), Liberia, Mexico, Netherlands Antilles, Nicaragua, Panama, Peru, Philippines, Saudi Arabia, Suriname, Taiwan, Trinidad Tobago, United States of America, Venezuela
6952301	Chicago, United States of America.

Related topics:

- Appendix F, “Parts listings”, on page 405

Appendix G. Customer-replaceable components

This topic describes the customer-replaceable components for the SAN File System engine and the master console.

Replacing storage engine components

Table 24. Model 1RX parts listing

Part number	Units	Description
18P6148	2	SAN File System Model 1RX
49P2167	2	• Power supply, AC 514 watt, hot-swap
01R0587	8	• Fan assembly, 80-mm hot-swap

Related topics:

- "Replacing a hot-swap fan" on page 123
- "Replacing a hot-swap power supply" on page 125
- "Replacing master console components"

Replacing master console components

Table 25. Master console parts listing

Part number	Units	Description
	1	Master console components
32P1032	1	• Display
28L3644	1	• Keyboard

Refer to the *IBM xSeries 345 Type 8670 Hardware Maintenance Manual and Troubleshooting Guide*, which was shipped with your master console, for replacement instructions.

Related topics:

- "Replacing storage engine components"

Appendix H. Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:
INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
MW9A/050
5600 Cottle Road
San Jose, CA 95193
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice and represent goals and objectives only.

Electronic emission notices

Related topics:

- "Federal Communications Commission (FCC) statement"
- "Australia and New Zealand Class A statement" on page 415
- "Industry Canada Class A emission compliance statement" on page 415
- "Chinese Class A warning statement" on page 417
- "European Union EMC Directive conformance statement" on page 415
- "Japanese Voluntary Control Council for Interference (VCCI) statement" on page 417
- "Taiwan electrical emission statement" on page 417
- "United Kingdom telecommunications safety requirement" on page 415

Federal Communications Commission (FCC) statement

Federal Communications Commission (FCC) Class A Statement:

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Note: 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 harmful 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 or her own expense.

Industry Canada Class A emission compliance statement

This Class A digital apparatus complies with Canadian ICES-003.

Avis de conformité à la réglementation d'Industrie Canada

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

Australia and New Zealand Class A statement

Attention: 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 adequate measures.

United Kingdom telecommunications safety requirement

Notice to Customers

This apparatus is approved under approval number NS/G/1234/J/100003 for indirect connection to public telecommunication systems in the United Kingdom.

European Union EMC Directive conformance statement

This product is in conformity with the protection requirements of EU Council Directive 89/336/EEC on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a nonrecommended modification of the product, including the fitting of non-IBM option cards.

This product has been tested and found to comply with the limits for Class A Information Technology Equipment according to CISPR 22/European Standard EN 55022. The limits for Class A equipment were derived for commercial and industrial environments to provide reasonable protection against interference with licensed communication equipment.

Attention: 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 adequate measures.

Germany Electromagnetic Compatibility Directive

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) vom 18. September 1998 (bzw. der EMC EG Richtlinie 89/336)

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Konformitätserklärung nach Paragraph 5 des EMVG ist die: IBM Deutschland Informationssysteme GmbH 70548 Stuttgart.

Informationen in Hinsicht EMVG Paragraph 4 Abs. (1) 4:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55022 Klasse A.

EN 55022 Klasse A Geräte müssen mit folgendem Warnhinweis versehen werden: “Warnung: dies ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funkstörungen verursachen; in diesem Fall kann vom Betreiber erlangt werden, angemessene Maßnahmen durchzuführen und dafür Ufzukommen.”

Anmerkung: Um die Einhaltung des EMVG sicherzustellen, sind die Geräte wie in den IBM Handbüchern angegeben zu installieren und zu betreiben.

Chinese Class A warning statement

Attention: 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 adequate measures.

中华人民共和国“A类”警告声明

声 明

此为A级产品，在生活环境中，该产品可能会造成无线电干扰。在这种情况下，可能需要用户对其干扰采取切实可行的措施。

Taiwan electrical emission statement

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。

Japanese Voluntary Control Council for Interference (VCCI) statement

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。

End of life statement

This box is a purchased unit. Therefore, it is the sole responsibility of the purchaser to dispose of it in accordance with local laws and regulations at the time of disposal.

The unit contains recyclable materials. The materials should be recycled where facilities are available and according to local regulations. In some areas, IBM may provide a product take-back program that ensure correct handling of the product. Contact your IBM representative for more information.

Safety information

The following sections describe the safety and environmental items you must consider before working with a Model 1RX.

Related topics:

- "Basic safety information (multilingual translations)"
- "General safety" on page 422
- "Grounding (earthing) requirements" on page 423
- "Handling electrostatic discharge-sensitive devices" on page 423
- "Handling static-sensitive devices" on page 424
- "Safety inspection guide" on page 424
- "System reliability considerations" on page 425
- "Working inside the engine with the power on" on page 426
- "Before you begin replacing components" on page 426

Basic safety information (multilingual translations)

DANGER

Before you begin to install this product, read the safety information in *Caution: Safety Information–Read This First*, **SD21-0030. This booklet describes safe procedures for cabling and plugging in electrical equipment.**

Gevarr: Voodrat u begint met de installatie van dit produkt, moet u eerst de veiligheidsinstructies lezen in de brochure *PAS OP! Veiligheidsinstructies–Lees dit eerst*, SD21-0030. Hierin wordt beschreven hoe u electrische apparatuur op een veilige manier moet bekabelen en aansluiten

Danger: Avant de procéder à l'nstallation de ce produit, lisez d'abord les consignes de sécurité dans la brochure *ATTENTION: Consignes de sécurité–A lire au préalable*, SD21-0030. Cette brochure décrit les procédures pour câbler et connecter les appareils électriques en toute sécurité.

Perigo: Antes de começar a instaler deste produto, leia as informações de segurança contidas em *Cuidado: Informações Sobre Segurança–Leia Primeiro*, SD21-0030. Esse folheto descreve procedimentos de segurança para a instalação de cabos e conexões em equipamentos elétricos.

危險：安裝本產品之前， 請先閱讀 "Caution: Safety Information--Read This First" SD21-0030 手冊中所提供的安全注意事項。 這本手冊將會說明使用電器設備的纜線及電源的安全程序。

Opasnost: Prije nego sto počnete sa instalacijom produkta, pročitajte naputak o pravilima o sigurnom rukovanju u Upozorenje: Pravila o sigurnom rukovanju - Prvo pročitaj ovo, SD21-0030. Ovaj privitak opisuje sigurnosne postupke za priključrivanje kabela i priključivanje na električno napajanje.

Upozornění: než zahájíte instalaci tohoto produktu, přečtěte si nejprve bezpečnostní informace v pokynech „Bezpečnostní informace“ č. 21-0030. Tato brožurka popisuje bezpečnostní opatření pro kabeláž a zapojení elektrického zařízení.

Fare! Før du installerer dette produkt, skal du læse sikkerhedsforskrifterne i *NB: Sikkerhedsforskrifter – Læs dette først* SD21-0030. Vejledningen beskriver den fremgangsmåde, du skal bruge ved tilslutning af kabler og udstyr.

Gevarr: Voordat u begint met het installeren van dit produkt, dient u eerst de veiligheidsrichtlijnen te lezen die zijn vermeld in de publikatie *Caution: Safety Information - Read This First*, SD21-0030. In dit boekje vindt u veilige procedures voor het aansluiten van elektrische appratuur.

VARRA: Ennen kuin aloitat tämän tuotteen asennuksen, lue julkaisussa *Varoitus: Turvaohjeet–Lue tämä ensin*, SD21-0030, olevat turvaohjeet. Tässä kirjasessa on ohjeet siitä, mitensähkölaitteet kaapeloidaan ja kytketään turvallisesti.

Danger : Avant d'installer le présent produit, consultez le livret *Attention : Informations pour la sécurité–Lisez-moi d'abord*, SD21-0030, qui décrit les procédures à respecter pour effectuer les opérations de câblage et brancher les équipments électriques en toute sécurité.

Vorsicht: Bevor mit der Installation des Produktes begonnen wird, die Sicherheitshinweise in *Achtung: Sicherheitsinformationen–Bitte zuerst lesen*. IBM Form SD21-0030. Diese Veröffentilchung beschreibt die Sicherheitsvorkehrungen für das Verkabien und Anschließen elektrischer Geräte.

Κίνδυνος: Πριν ξεκινήσετε την εγκατάσταση αυτού του προϊόντος, διαβάστε τις πληροφορίες ασφάλειας στο φυλλάδιο *Caution: Safety Information-Read this first*, SD21-0030. Στο φυλλάδιο αυτό περιγράφονται οι ασφαλείς διαδικασίες για την καλωδίωση των ηλεκτρικών συσκευών και τη σύνδεσή τους στην πρίζα.

Vigyázat: Mielôtt megkezdi a berendezés üzembe helyezését, olvassa el a *Caution: Safety Information–Read This First*, SD21-0030 könyvecskében leírt biztonsági információkat. Ez a könyv leírja, miyen biztonsági intézkedéseket kell megtenni az elektromos berendezés huzalozásakor illetve csatlakoztatásakor.

Pericolo: prima di iniziare l'installazione di questo prodotto, leggere le informazioni relatie alla sicurezza riportate nell'opuscolo *Attenzione: Informazioni di sicurezza–Prime informazioni da leggere* in cui sono descritte le procedure per il cablaggio ed il collegamento di apparecchiature elettriche.

危険： 導入作業を開始する前に、安全に関する
小冊子SD21-0030 の「最初にお読みください」
(Read This First)の項をお読みください。
この小冊子は、電気機器の安全な配線と接続の
手順について説明しています。

위험: 이 제품을 설치하기 전에 반드시
"주의: 안전 정보-시작하기 전에"
(SD21-0030) 에 있는 안전 정보를
읽으십시오.

ОПАСНОСТ
Пред да почнете да го инсталирате овој продукт, прочитајте
ја информацијата за безбедност:
"Предупредување: Информација за безбедност: Прочитајте го
прво ова", SD21-0030.
Оваа брошура опишува безбедносни процедури за каблирање
и вклучување на електрична опрема.

Fare: Før du begynner å installere dette produktet, må du lese sikkerhetsinformasjonen i *Advarsel: Sikkerhetsinformasjon – Les dette forst*, SD21-0030 som beskriver sikkerhetsrutinene for kabling og tilkobling av elektrisk utstyr.

Uwaga:
Przed rozpoczęciem instalacji produktu należy zapoznać się z instrukcją:
"Caution: Safety Information - Read This First", SD21-0030.
Zawiera ona warunki bezpieczeństwa przy podłączaniu do sieci elektrycznej
i eksploatacji.

Perigo: Antes de iniciar a instalação deste produto, leia as informações de segurança *Cuidado: Informações de Segurança–Leia Primeiro*, SD21-0030. Este documento descreve como efectuar, de um modo seguro, as ligações eléctricas dos equipamentos.

ОСТОРОЖНО: Прежде чем инсталлировать этот
продукт, прочтите Инструкцию по технике безо-
пасности в документе "Внимание: Инструкция по
технике безопасности -- Прочесть в первую очередь",
SD21-0030. В этой брошюре описаны безопас-
ные способы каблирования и подключения элект-
рического оборудования.

Nebezpečenstvo: Pred inštaláciou výrobku si prečítajte bezpečnosté predpisy v
Výstraha: Bezpeč osté predpisy - Prečítaj ako prvé, SD21 0030. V tejto brožúrke sú opísané bezpečnosté postupy pre pripojenie elektrických zariadení.

Pozor: Preden zaènete z instalacijo tega produkta preberite poglavje: 'Opozorilo: Informacije o varnem rokovanju-preberi pred uporabo," SD21-0030. To poglavje opisuje pravilne postopke za kabliranje,

Peligro: Antes de empezar a instalar este producto, lea la información de seguridad en *Atención: Información de Seguridad–Lea Esto Primero*, SD21-0030. Este documento describe los procedimientos de seguridad para cablear y enchufar equipos eléctricos.

Varning — livsfara: Innan du börjar installera den här produkten bör du läsa säkerhetsinformationen i dokumentet *Varning: Säkerhetsföreskrifter – Läs detta först*, SD21-0030. Där beskrivs hur du på ett säkert satt ansluter elektrisk utrustning.

危險：

開始安裝此產品之前，請先閱讀安全資訊。

注意：

請先閱讀 - 安全資訊 SD21-0030

此冊子說明插接電器設備之電纜線的安全程序。

General safety

Follow these rules to ensure general safety:

- Observe good housekeeping in the area of the machines during and after maintenance.
- When lifting any heavy object:
  1. Ensure that you can stand safely without slipping.
  2. Distribute the weight of the object equally between your feet.
  3. Use a slow lifting force. Never move suddenly or twist when you attempt to lift.

4. Lift by standing or by pushing up with your leg muscles; this action removes the strain from the muscles in your back. *Do not attempt to lift any objects that weigh more than 16 kg (35 lb.) or objects that you think are too heavy for you.*

- Do not perform any action that causes hazards to the customer, or that makes the equipment unsafe.
- Before you start the machine, ensure that other service representatives and the customer's personnel are not in a hazardous position.
- Place removed covers and other parts in a safe place, away from all personnel, while you are servicing the machine.
- Keep your tool case away from walk areas so that other people will not trip over it.
- Do not wear loose clothing that can be trapped in the moving parts of a machine. Ensure that your sleeves are fastened or rolled up above your elbows. If your hair is long, fasten it.
- Insert the ends of your necktie or scarf inside clothing or fasten it with a nonconductive clip, approximately 8 centimeters (3 inches) from the end.
- Do not wear jewelry, chains, metal-frame eyeglasses, or metal fasteners for your clothing.

  Remember: Metal objects are good electrical conductors.
- Wear safety glasses when you are: hammering, drilling soldering, cutting wire, attaching springs, using solvents, or working in any other conditions that might be hazardous to your eyes.
- After service, reinstall all safety shields, guards, labels, and ground wires. Replace any safety device that is worn or defective.
- Reinstall all covers correctly before returning the machine to the customer.

Grounding (earthing) requirements

Electrical grounding (earthing) of the computer is required for operator safety and correct system function. Proper grounding of the electrical outlet can be verified by a certified electrician.

Handling electrostatic discharge-sensitive devices

Any computer part containing transistors or integrated circuits (ICs) should be considered sensitive to electrostatic discharge (ESD). ESD damage can occur when there is a difference in charge between objects. Protect against ESD damage by equalizing the charge so that the machine, the part, the work mat, and the person handling the part are all at the same charge.

Notes:

1. Use product-specific ESD procedures when they exceed the requirements noted here.
2. Make sure that the ESD protective devices you use have been certified (ISO 9000) as fully effective.

When handling ESD-sensitive parts:

- Keep the parts in protective packages until they are inserted into the product.
- Avoid contact with other people.
- Wear a grounded wrist strap against your skin to eliminate static on your body.
- Prevent the part from touching your clothing. Most clothing is insulative and retains a charge even when you are wearing a wrist strap.

- Use the black side of a grounded work mat to provide a static-free work surface. The mat is especially useful when handling ESD-sensitive devices.
- Select a grounding system, such as those listed below, to provide protection that meets the specific service requirement.

  Note: The use of a grounding system is desirable but not required to protect against ESD damage.

  – Attach the ESD ground clip to any frame ground, ground braid, or green-wire ground.
  – Use an ESD common ground or reference point when working on a double-insulated or battery-operated system. You can use coaxial or connector-outside shells on these systems.
  – Use the round ground-prong of the AC plug on AC-operated computers.

Handling static-sensitive devices

Attention: Static electricity can damage electronic devices and your engine. To avoid damage, keep static-sensitive devices in their static-protective packages until you are ready to install them.

To reduce the possibility of electrostatic discharge, observe the following precautions:

- Limit your movement. Movement can cause static electricity to build up around you.
- Handle the device carefully, holding it by its edges or its frame.
- Do not touch solder joints, pins, or exposed printed circuitry.
- Do not leave the device where others can handle and possibly damage the device.
- While the device is still in its static-protective package, touch it to an unpainted metal part of the server for at least 2 seconds. (This drains static electricity from the package and from your body.)
- Remove the device from its package and install it directly into the server without setting it down. If it is necessary to set the device down, place it in its static-protective package. Do not place the device on the engine cover or on a metal table.
- Take additional care when handling devices during cold weather because heating reduces indoor humidity and increases static electricity.

Related topics:

- "Handling electrostatic discharge-sensitive devices" on page 423

Safety inspection guide

The intent of this inspection guide is to assist you in identifying potentially unsafe conditions on these products. Each machine, as it was designed and built, had required safety items installed to protect users and service personnel from injury. This guide addresses only those items. However, good judgment should be used to identify potential safety hazards due to attachment of non-IBM features or options not covered by this inspection guide.

If any unsafe conditions are present, you must determine how serious the apparent hazard could be and whether you can continue without first correcting the problem.

Consider these conditions and the safety hazards they present:
- Electrical hazards, especially primary power (primary voltage on the frame can cause serious or fatal electrical shock)
- Explosive hazards, such as a damaged CRT face or bulging capacitor
- Mechanical hazards, such as loose or missing hardware

The guide consists of a series of steps presented in a checklist. Begin the checks with the power off and the power cord disconnected.

Checklist:

1. Check exterior covers for damage (loose, broken, or sharp edges).
2. Power OFF the engine. Disconnect the power cord.
3. Check the power cord for:
   a. A third-wire ground connector in good condition. Use a meter to measure third-wire ground continuity for 0.1 ohm or less between the external ground pin and frame ground.
   b. The power cord should be the appropriate type as specified in the parts listings.
   c. Insulation must not be frayed or worn.
4. Remove the cover.
5. Check for any obvious non-IBM alterations. Use good judgment as to the safety of any non-IBM alterations.
6. Check inside the unit for any obvious unsafe conditions, such as metal filings, contamination, water or other liquids, or signs of fire or smoke damage.
7. Check for worn, frayed, or pinched cables.
8. Check that the power-supply cover fasteners (screws or rivets) have not been removed or tampered with.

Related topics:
- "General safety" on page 422
- "Safety information" on page 418

System reliability considerations

This section is for the IBM service representative.

To help ensure proper cooling and system reliability, make sure that:
- Each of the drive bays has either a drive or a filler panel installed.
- Each of the power-supply bays has a power supply installed.
- For rack configurations, make sure that space is available around the engine to enable the engine cooling system to work properly. See the documentation that comes with the rack for additional information.
- The engine cover is in place during normal operation.
- The air-baffle cover over the microprocessors remains closed during normal operation.
- The air baffle is installed between the fans and the power supply.
- A removed hot-swap drive is replaced within 10 minutes of removal.
- Cables for optional adapters are routed according to the instructions provided with the adapters.
- A failed fan is replaced within 48 hours.

- The engine is powered off and the power cords are disconnected before you remove an air baffle.
- The air baffle is always installed in the engine except when you are installing or removing the components that are located under the air baffle.

Working inside the engine with the power on

The Model 1RX supports hot-plug, hot-add, and hot-swap devices and is designed to operate safely while powered on with the cover removed. Follow these guidelines when you work inside a engine that is powered on:

- Avoid loose-fitting clothing on your forearms. Button long-sleeved shirts before working inside the engine; do not wear cuff links while you are working inside the engine.
- Do not allow your necktie or scarf to hang inside the engine.
- Remove jewelry, such as bracelets, necklaces, rings, and loose-fitting wristwatches.
- Remove items from your shirt pocket (such as pens or pencils) that could fall into the engine as you lean over it.
- Avoid dropping any metallic objects, such as paper clips, hairpins, or screws, into the engine.

Note: The adapters that are preinstalled in the Model 1RX are not hot pluggable. You must power off the engine and disconnect the power cords to replace an adapter.

Before you begin replacing components

Before you begin to install options in the Model 1RX, read the following information:

- Become familiar with the safety and handling guidelines and read the safety statements.
- You do not need to power off the engine to install or replace hot-swap power supplies, hot-swap hard disk drives, or hot-swap fans.
- The orange color on components and labels in your engine identifies hot-swap or hot-plug components. You can install or remove hot-swap and hot-plug components while the engine is running, provided that your engine is configured to support this function.
- The blue color on components and labels identifies touch points where you can grip a component, move a latch, and so on.
- Make sure that you have an adequate number of properly grounded electrical outlets for the engine, monitor, and any other options that you intend to install.
- Back up all important data before you make changes to disk drives.

Related topics:

- "Handling static-sensitive devices" on page 424
- "Safety information" on page 418
- "System reliability considerations" on page 425
- "Working inside the engine with the power on"

Trademarks

The following terms are trademarks of International Business Machines Corporation or Tivoli Systems Inc. in the United States or other countries or both:

AIX	IBM	Tivoli
DB2	IBM logo	TotalStorage
Enterprise Storage Server	SecureWay	WebSphere
FlashCopy	StorageTank	xSeries

Java and all Java-based trademarks are registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks or registered trademarks of Microsoft Corporation.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, and service names may be trademarks or service marks of others.

Glossary

This glossary includes terms and definitions from:

- *The American National Standard Dictionary for Information Systems*, ANSI X3.172-1990, copyright 1990 by the American National Standards Institute (ANSI). Copies can be purchased from the American National Standards Institute, 1430 Broadway, New York, New York 10018. Definitions are identified by the symbol (A) after the definition.
- The *ANSI/EIA Standard - 440A: Fiber Optic Terminology*, copyright 1989 by the Electronics Industries Association (EIA). Copies can be purchased from the Electronics Industries Association, 2001 Pennsylvania Avenue N.W., Washington, D.C. 20006. Definitions are identified by the symbol (E) after the definition.
- The *Information Technology Vocabulary*, developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC1/SC1). Definitions of published parts of this vocabulary are identified by the symbol (I) after the definition; definitions taken from draft international standards, committee drafts, and working papers being developed by ISO/IEC JTC1/SC1 are identified by the symbol (T) after the definition, indicating that final agreement has not yet been reached among the participating National Bodies of SC1.
- The *Storage Networking Dictionary*, available online at the Storage Networking Industry Association (SNIA) Web site:

  www.snia.org/education/dictionary/
- The Distributed Management Task Force (www.dmtf.org), copyright 2003 by the Distributed Management Task Force, Inc., 225 SE Main Street Portland, OR 97214. Definitions derived from this book have the symbol (D) after the definition.

This glossary uses the following cross-reference forms:

See This refers the reader to one of two kinds of related information:

- A term that is the expanded form of an abbreviation or acronym. This expanded form of the term contains the full definition.
- A synonym or more preferred term

See also This refers the reader to one or more related terms.

ACLI. See *Administrative command-line interface (ACLI)*.

Administrative command-line interface (ACLI). A command-line interface used to administer all aspects of the SAN File System. The ACLI runs on all engines that host Metadata servers and the Administrative server.

administrative log. A log that maintains a history of messages created by the Administrative server.

Administrative server. For SAN File System, a set of servlets running within a customized instance of WebSphere Application Server that handles all SAN File System administrative requests from the SAN File System console. See also *SAN File System console*.

alert. A message or other indication that identifies a problem or an impending problem.

audit log. A log that maintains the history of all commands issued by any administrator for all Metadata servers in the cluster.

CIM. See *Common Information Model*.

CIM client application. A storage management program that initiates CIM requests to the Administrative agent for the device.

CIM namespace. The scope within which a CIM schema applies.

CIM object manager (CIMOM). The common conceptual framework for data management that receives, validates, and authenticates the CIM requests from the client application and then directs the requests to the appropriate component or device provider.

CIMOM. See *CIM object manager*.

client. For SAN File System, a client is a system that can access the SAN File System. These clients act as servers to a broader clientele, providing Network File System or Common Internet File System access to the global namespace or hosting applications (such as database servers or Web-hosting services that use multiple servers).

class. The definition of an object within a specific hierarchy. An object class can have properties and methods and serve as the target of an association.

CLI. See *Administrative command-line interface*.

client state manager (CSM). A component of the client kernel that provides protocol support for the client.

cluster. A group of engines that is managed as a set and presents a single point of control for configuration and service activity.

cluster log. A log that maintains a history of messages created by all Metadata servers in the cluster.

cluster state. A status condition of the cluster. Cluster states can be inactive (Not running or Forming), active (Online, Offline, Partly quiescent, or Fully quiescent) or unknown. See also *Forming, Fully quiescent, Not running, Offline, Online,* and *Partly quiescent*.

Common Information Model (CIM). A set of standards from the Distributed Management Task Force Inc. (DMTF). CIM provides a conceptual framework for storage management and an open approach to the design and implementation of storage systems, applications, databases, networks, and devices.

coordinated universal time (UTC). The time scale, based on the System International (SI) second, as defined and recommended by the Comitb Consultatif International de la Radio (CCIR) and maintained (using an atomic clock) by the Bureau International des Poids et Mesures (BIPM).

CSM. See *client state manager*.

default user storage pool. A storage pool that stores file data that SAN File System has not assigned (using the active policy) to a user storage pool, as well as file data that is assigned directly to this storage pool. There is only one default user storage pool; however, you can assign any user storage pool as the default storage pool. See also *user storage pool*

engine. The hardware unit that hosts the software for the Metadata server.

event log. (1) A log that maintains a history of event messages issued by all Metadata servers in the cluster. (2) IBM Term: A log that contains information about events for a particular system or group, for a particular metric, or for all the events that are associated with a specific monitor.

file-placement rule. A rule that controls in what pool SAN File System places files in the global namespace. See also *rule* and *global namespace*.

fileset. A hierarchical grouping of files managed as a unit for balancing workload across a cluster.

FlashCopy image. A space-efficient image of the contents of part of the SAN File System at a particular moment.

Forming. A status condition where the cluster has a master and is in the process of forming. This state is always the initial one whenever a cluster is newly formed.

Fully quiescent. A status condition that cuts off all client communication with the cluster.

global namespace. A single file system that provides complete, shared access to both Windows and UNIX clients in the same environment.

ID. See *identifier*.

identifier (ID). A sequence of bits or characters that identifies a user, program, device, or system to another user, program, device, or system.

Initializing. A status condition during which a Metadata server or the entire cluster is set up for the first time.

key. A property that is used to provide a unique identifier for an instance of a class. Key properties are marked with the Key qualifier. (D)

lease. The amount of time that a client can hold a lock.

lock. A restriction that allows clients to have exclusive access to files. Types of locks include *data locks*, *session locks*, and *range locks*.

logical unit (LU). In open systems, a logical disk drive.

logical unit number (LUN). In the small computer system interface (SCSI) protocol, a unique number used on a SCSI bus to enable it to differentiate between up to sixteen separate devices per SCSI ID address, each of which is a logical unit.

LU. See *logical unit*.

LUN. See *logical unit number*.

managed object format (MOF). A compiled language for defining classes and instances. A MOF compiler offers a textual means of adding data to the CIM Object Manager repository. MOF eliminates the need to write code, thus providing a simple and fast technique for modifying the CIM Object Manager repository. (D)

master Metadata server. In SAN File System, the Metadata server in a cluster that is responsible for physical-space allocation.

metadata. Data that describes the characteristics of stored data; descriptive data.

Metadata server. In SAN File System, a server that offloads the metadata processing from the data-storage environment to improve SAN performance. An instance

of the SAN File System runs on each engine, and together the Metadata servers form a cluster. See also *cluster*.

method. A way to implement a function on a class.

MOF. See *managed object format*.

Not running. A status condition where one or more servers in the cluster are not added and therefore the cluster cannot perform any functions.

object name. An object that consists of a CIM namespace path and a model path. The namespace path provides access to the CIM implementation managed by the CIM agent, and the model path provides navigation within the implementation.

Offline. A status condition during which clients are not being serviced and the cluster is responding only to administrative requests.

Online. A status condition that indicates the normal operational state for the cluster.

Partly Quiesced. A state in which the cluster or server is in a "quiet" client communications mode to allow other operations to occur.

Partly quiescent. A status condition that allows existing metadata activity and client communication to continue on the cluster, but prohibits new communication.

policy. A list of file-placement rules that define characteristics and placement of files. Several policies can be defined within the configuration, but only one policy is active at one time. See also *file-placement rule* and *service-class rule*.

pool. See *storage pool*.

property. An attribute that is used to characterize instances of a class.

qualifier. A value that provides additional information about a class, association, indication, method, method parameter, instance, property, or reference.

quota. A limit on the amount of disk space a user can use.

rule. The lines within a policy that specify which actions will occur when certain conditions are met. Conditions include attributes about an object (file name, type or extension, dates, owner, and groups) and the fileset name associated with the object.

SAN File System console. A Web user interface used to monitor and control the SAN File System remotely by using any standard Web browser.

schema. A group of object classes defined for and applicable to a single namespace. Within the CIM agent, the supported schemas are loaded through the managed object format (MOF) compiler.

security log. A log that maintains a history of Administrative server login activity.

service location protocol. A directory service that the CIM client application calls to locate the CIMOM.

Shutdown. A status condition that describes when the cluster is shut down as intended.

SLP. See *service location protocol*.

Starting. A status condition when a Metadata server is starting as designed but is not ready to accept connections from clients.

storage pool. A named set of storage volumes that is the destination for storing client data.

symbolic link. A type of file that contains the path name of and acts as a pointer to another file or directory.

system storage pool. A storage pool that contains the system metadata (system and file attributes, configuration information, and Metadata server state) that is accessible to all Metadata servers in the cluster. There is only one system storage pool. See also *Metadata server*.

user storage pool. An optional storage pool that contains blocks of data that compose the files that are created by SAN File System clients. See also *storage pool* and *default user storage pool*.

volume. A labeled logical unit, which can be a physical device or a logical device. For SAN File System, there is a one to one relationship between volumes and LUNs. See also *logical unit number*.

UTC. See *coordinated universal time*

Index

A

About the Maintenance and Problem Determination Guide vii
AC power LED 366
accessibility
  disability 153
  keyboard 153
  shortcut keys 153
activatevol 163
activating
  volume 163
activating a policy 296
adapter 111
  retention bracket 363
  support bracket 363
adapter considerations 111
adapter, PCI bus, identification 111
adding
  SNMP manager 164
  volume to a storage pool 240
adding a Metadata server to the cluster 322
addserver 322
addsnmpmgr 164
administrative
  log 37
  server 8
Administrative agent
  starting 343
  stopping 351
Administrative CLI (ACLI)
  accessing 157
Administrative server
  troubleshooting 67
administrative user
  creating 246
administrative users
  listing 187, 227
Advanced System Management ports 374
air-baffle assembly, illustration 363
AIX client
  commands 297
AIX, client for
  controlling disk access 306
  creating 303
  destroying 303
  displaying status for 313
  loading the file-system driver 309
  mounting the global namespace 311
  scanning SAN File System for new and removed volumes 306
  statistics 347
assigning
  filesets to a new Metadata server 260
attachcontainer 166
attaching a fileset 166
audit log
  clearing 180
Audit log 38
auxiliary-device port 373

B

backing up system metadata 65
backplane, replacing 120
battery
  connector 371, 372
battery, replacing 117
beep symptoms 377
bezel, installing 144
BIOS code, recovering 101
boot option, adapter 111
builddrscript 168
button
  CD-eject 365
  Diskette-eject 365
  Power-control 365
  Reset 365

C

cable
  connectors 367, 368
  SCSI 375
cabling
  Ethernet 373
  external 375
  system-planar external connectors 367
  system-planar internal connectors 368
capturing first failure data 337
catlog 169
catpolicy 171
CD-eject button 365
CD-ROM drive
  Activity LED 365
CD-ROM drive, replacing 119
CD, publications viii
certificates, replacing expired 73
chadvclusterconfig 323
chclusterconfig 172
chcontainer 174
chpool 176
chvol 178
CIMOM certificate 73
Class A electronic emission notice 414
Class A statement
  Australia 415
  Canada 415
  China 417
  New Zealand 415
clearing
  audit log 180
  cluster log 180
clearlog 180
client
  commands 297
  creating 302, 303
  debugging for AIX 345
  destroying 303
  displaying status for 313
client *(continued)*
  listing active 194
  loading the file-system driver 302, 309
  mounting the global namespace 302, 311
  removing 301
  setting up 302
  unmounting the global namespace 314
Client CLI
  accessing 157
client diagnostic tools 41, 155
client logging and tracing 42, 44, 85
clients
  supported 10
  troubleshooting performance 80
  UNIX-based 12
  Windows-based 13
cluster
  displaying status 278
  log 40
  modifying advanced settings 323
  modifying settings 172
  starting 272
  stopping 290
  upgrading 295
cluster log
  clearing 180
cluster, adding a Metadata server to 322
cluster, dropping a Metadata server from 326
collectdiag 181
collecting diagnostic data 181
command
  activatevol 163
  addserver 322
  addsnmpmgr 164
  attachcontainer 166
  builddrscript 168
  catlog 169
  catpolicy 171
  chadvclusterconfig 323
  chclusterconfig 172
  chcontainer 174
  chpool 176
  chvol 178
  clearlog 180
  collectdiag 181
  delog 34
  detachcontainer 183
  disableConsoleTrace 325
  dropserver 326
  elog 34
  enableConsoleTrace 327
  exit 185
  help 186
  legacy 328
  lsadmuser 187
  lsautorestart 190
  lsclient 194

command *(continued)*
lscontainer 198
lsdrfile 202
lsengine 204
lsimage 207
lslun 210
lspolicy 213
lspool 216
lsproc 219
lsserver 221
lssnmpmgr 225
lstrapsetting 226
lsvol 227
migratedata 298, 316
mkcontainer 230
mkdrfile 233
mkimage 234
mkpolicy 236
mkpool 238
mktruststore 335
mkvol 240
pmf 337
quiescecluster 242
quit 244
reportvolfiles 245
resetadmuser 246
restartengine 247
resumecluster 248
reverttoimage 249
rmcontainer 251
rmdrfile 252
rmimage 253
rmpolicy 255
rmpool 256
rmsnmpmgr 257
rmstclient 301
rmvol 258
setcontainerserver 260
setdefaultpool 262
setmaster 263
setoutput 266
settrap 268
setupstclient 302
SHOW 328
startautorestart 270
startcluster 272
startConsole 344
startengine 273
startmetadatacheck 274
startserver 276
statcluster 278
statcontainer 281
statengine 282
statserver 286
stfsclient 303
stfsdebug 345
stfsdisk 306
stfsdriver 309
stfsmount 311
stfsstat 347
stfsstatus 313
stfsumount 314
stopautorestart 289
stopcluster 290
stopConsole 352
stopengine 291
stopmetadatacheck 292
command *(continued)*
stopserver 293
suspendvol 294
tank extractbootrecord 353
tank lscluster 355
tank lsdisklabel 358
tank lsversion 357
tank resetcluster 360
tank resetversion 362
TRACE 328
usepolicy 296
vpdlog 35
command-line interface
troubleshooting user access 70
commands 319
administrative 158
AIX client 297
client 297
troubleshooting user access 71
Windows client 315
components
major, Model 1RX 363
SAN File System 3
system planar 371, 372
Configuration/Setup Utility 106
configuring LDAP 74
connectors
adapter 371, 372
battery 371, 372
cable 368
Ethernet 373
external port 367
internal cable 368
memory 371, 372
microprocessor 371, 372
port 367
SCSI 374
system planar 371, 372
system-planar switch block 369
Universal Serial Bus (USB) 375
VRM 371, 372
console
troubleshooting user access 68
controlling disk access for a client for AIX 306
controls 365
cover, removing 110
creating
client 302
client for AIX 303
disaster-recovery dump files 233
disaster-recovery script files 168
filesets 230
FlashCopy image 234
policies 236
storage pools 238
truststore 335
customer-replaceable components 411

D

data, collecting diagnostic 181
DC power LED 366
debugging client for AIX 345
deleting
disaster-recovery dump files 252
FlashCopy image 253
deleting *(continued)*
storage pool 256
delog command 34
destroying
client for AIX 303
detachcontainer 183
detaching a fileset 183
details
active and inactive policies 213
administrative users 187
available LUNs 210
client 194
cluster 278
disaster-recovery files 202
engine 282
engines 204
fileset 198
FlashCopy image 207
Metadata server 286
diagnostic tools 29
engine 29
diags command 33
DIMM 127
disableConsoleTrace 325
disabling
automatic restart of the Metadata server 289
SAN File System console Web server 325
disaster recovery 145
disaster recovery commands 353, 355, 357, 358, 360, 362
disaster-recovery dump files
creating 233
deleting 252
disaster-recovery files
listing 202
disaster-recovery script files
creating 168
disk drive recovery
master console 84
diskette drive
Activity LED 365
diskette drive, replacing 121
Diskette-eject button 365
displaying
automatic restart of the Metadata server 190
cluster status 278
engine status 282
log-file content 169
long-running processes 219
Metadata server status 286
policy rules 171
SNMP managers 225
statistics for a client for AIX 347
statistics for a client for Windows 347
status for a client for AIX 313
drive, hot-swap, installing 122
dropping a Metadata server from the cluster 326
dropserver 326
dump capability
Linux 41, 43
SAN File System 41
Windows 2000 45

E

earthing (grounding) requirements 423
electronic emission
  Class A notice 414
  notices 414
electrostatic discharge-sensitive devices, handling 423
elog command 34
enableConsoleTrace 327
enabling
  automatic restart of the Metadata server 270
end of life statement 417
ending a tanktool session 185, 244
engine
  adding or replacing components 109
  definition 5
  diagnostic tools 29
  displaying status 282
  listing 204
  powering off 291
  restarting 247
  starting 273
engine diagnostics, running 104
error codes, POST 389
error messages
  bus fault 380
  fan 386
  hard disk drive 381
  power 397
  temperature 399
ESD (electrostatic) devices, handling 423
Ethernet
  cabling 373
  connector 373
  pin-number assignment 373
Ethernet controller
  troubleshooting 90
Ethernet LEDs
  activity 366
  link status 366
European Union EMC Directive conformance statement 415
exit 185

F

fan
  error messages 386
fan, replacing 123
FCC Class A notice 414
fibre-channel cable replacement
  master console 85
file-system drive
  version of 313
files
  listing, in a volume 245
fileset
  attaching 166
  creating 230
  detaching 183
  listing 198
  modifying settings 174
  reassigning to a Metadata server 260
  removing 251
filler panel, hard disk drive bay 122
first failure data capture 281, 337
FlashCopy image
  creating 234
  deleting 253
  listing 207
  reverting to 249

G

GBIC replacement
  master console 85
Germany Electromagnetic Compatibility Directive 416
getting help 186
getting started 1
global namespace
  attaching a fileset to 166
grounding (earthing) requirements 423

H

handling static-sensitive devices 424
hard disk backplane, replacing 120
hard disk drive
  Activity LED 365
  Status LED 365
hard disk drive error messages 381
hardware
  features 7
  requirements 6
  specifications 6
hardware messages and codes 377
hardware vital product data 46
heat sink 131
help 186
  general 151
  online 152
  telephone 152
hot-swap and hot-plug devices
  drives 363
  fans 363
hot-swap drive, installing 122
hot-swap power supply, installing 125

I

IDE bus activity LED 365
indicators 365
Information panel 365
installing
  options 426
installing hot-swap fan 123
installing hot-swap power supply 125
installing memory module 127
installing non-hot-plug adapter 111
installing VRM 131
installing, hot-swap drive 122
installing, microprocessor 131
integrated
  Ethernet 373

J

Japanese Voluntary Control Council for Interference (VCCI) statement 417

K

keyboard port 374

L

LDAP certificate 73
LDAP configuration 74
LEDs
  AC power 366
  CD-ROM drive activity 365
  DC power 366
  Diskette-drive activity 365
  Ethernet activity 366
  Ethernet-link status 366
  Front view 365
  Hard disk drive activity 365
  Hard disk drive status 365
  IDE bus activity 365
  Power-on 365
  rear view 366
  SCSI bus activity 365
  System error 365
  system planar 369
  System-identification 365
legacy 328
legacy CLI commands, running 328
light path diagnostics
  overview 30, 31
  panel LEDs 387
limited warranty viii
listing
  active and inactive policies 213
  active clients 194
  administrative users 187, 227
  automatic restart of the Metadata server 190
  available LUNs 210
  disaster-recovery files 202
  engines 204
  files in a volume 245
  fileset 198
  FlashCopy image 207
  long-running processes 219
  Metadata server 221
  SNMP managers 225
  storage pool 216
  trap settings 226
loading the client file-system driver 302, 309
local area network
  Ethernet 373
local network
  troubleshooting 59
log
  administrative 37
  cluster 40
  security 39
log file
  displaying the content of 169
logs 36
  clearing 180
  engine 32
long-running processes
  listing 219
lsadmuser 187
lsautorestart 190

lsclient 194
lscontainer 198
lsdrfile 202
lsengine 204
lsimage 207
lslun 210
lspolicy 213
lspool 216
lsproc 219
lsserver 221
lssnmpmgr 225
lstrapsetting 226
lsvol 227
LUNs
  listing available 210
  troubleshooting 77

M

master console
  recovering disk drive 84
  replacing fibre-channel cable 85
Master console 16
master console, replacing components 411
memory module, installing 127
memory module, order of installation 127
messages
  voltage-related 402
messages and codes 377
metadata
  backing up system 65
  server 8
  starting the checker 274
  stopping the checker 292
Metadata server
  disabling automatic restart 289
  displaying status 286
  enabling automatic restart 190, 270
  listing 221
  setting the master 263
  starting 276
  starting all in the cluster 272
  stopping 293
  stopping all in the cluster 290
  troubleshooting 56
Metadata server, adding to the cluster 322
Metadata server, dropping from the cluster 326
microprocessor, installing 131
migratedata 298, 316
migrating data to SAN File System 298, 316
mkcontainer 230
mkdrfile 233
mkimage 234
mkpolicy 236
mkpool 238
mktruststore 335
mkvol 240
Model 1RX
  components 363
  components, illustration 364
  front view 365
  parts listing 405
Model 1RX *(continued)*
  rear view 366
modifying
  advanced cluster settings 323
  cluster settings 172
  fileset settings 174
  pool settings 176
  volume settings 178
mounting the global namespace on a client for AIX 311
mounting the global namespace on the client 302

N

navigating by keyboard 153
networks 373
notices 413
  electronic emission 414
  FCC, Class A 414
Notices viii

O

one-button data collection utility 45
operator panel card, replacing 136
options
  external
    cabling requirements 375
OS present diagnostics 33
output format
  setting 266
Overview 3

P

parts listings 405
PCI bus 111
pin-number assignments
  Ethernet 373
pmf 337
pointing device port 373
policies
  listing active and inactive 213
  removing 255
policy
  activating 296
  creating 236
policy rules
  displaying 171
pool
  creating 238
  modifying settings 176
  setting the default 262
port
  Advanced System Management 374
  auxiliary-device 373
  connectors 367
  input/output 372
  keyboard 374
  mouse 373
  parallel 375
  pointing device 373
  Universal Serial Bus (USB) 375
POST 30
POST error codes 389
power
  cords 407
  error messages 397
power supply, installing 125
Power-control button 365
Power-on LED 365
power-on password override switch 369
power-on self-test 30
power-supply cage, replacing 137
powering off an engine 291
problem determination, engine 101
processes
  listing long-running 219
publications viii
publications CD viii

Q

quiescecluster 242
  quiescing 242
  resuming 248
quiescing the cluster 242
quit 244

R

reassigning
  filesets to a new Metadata server 260
rebooting engine 247
recovering the BIOS code 101
release notes viii
reliability, system 425
remote access 17
removing
  client 301
  fileset 251
  policies 255
  SNMP manager 257
  volume 258
replacing
  master console components 411
  storage engine components 411
replacing battery 117
replacing CD-ROM drive 119
replacing diskette drive 121
replacing hard disk backplane 120
replacing hot-swap fan 123
replacing RSA II 139
replacing the diagnostics/operator panel card 136
replacing the power-supply cage 137
replacing the system planar and shuttle 142
reportvolfiles 245
reset button 365
resetadmuser 246
resetting
  administrative user 246
resolution procedures 72
restartengine 247
restarting
  Metadata server automatically 270, 289
restarting engine 247
resumecluster 248
resuming the cluster 248

reverting to a FlashCopy image 249
reverttoimage 249
rmcontainer 251
rmdrfile 252
rmimage 253
rmpolicy 255
rmpool 256
rmsnmpmgr 257
rmstclient 301
rmvol 258
RSA II
  logging on 106
  replacing the adapter 139
RSAII interface 35
running legacy CLI commands 328

S

safety
  general 422
safety information viii
  basic (multilingual translations) 418
  safety inspection guide 424
safety notices 418
safety notices, translated viii
SAN File System accessibility features 153
SAN File System console Web server
  disabling 325
  starting 344
  stopping 327, 352
SCSI
  cable
    external devices 375
    maximum length 375
  connector
    description 374
    location 368
  controller 374
  IDs 374
SCSI adapter boot option 111
SCSI bus activity LED 365
security
  log 39
server
  administrative 8
  metadata 8
server diagnostic tools 36
service
  phone 152
Service Alert 16
setcontainerserver 260
setdefaultpool 262
setmaster 263
setoutput 266
setting
  master Metadata server 263
  SNMP traps 268
setting output format 266
setting the default storage pool 262
setting up the clients 302, 313
settings
  listing, trap 226
settrap 268
setupstclient 302
SHOW command 328
Simple Network Management Protocol (SNMP)
  components 18
SNMP
  adding a manager 164
  listing managers 225
  removing a manager 257
  setting a trap 268
SNMP (Simple Network Management Protocol)
  components 18
software vital product data 47
startautorestart 270
startCimom 343
startcluster 272
startConsole 344
startengine 273
starting
  all Metadata servers in the cluster 272
  metadata checker 274
  SAN File System console Web server 344
starting an engine 273
starting an Metadata server 276
starting the Administrative agent 343
startmetadatacheck 274
startserver 276
statcluster 278
statcontainer 281
statengine 282
static electricity 424
static-sensitive devices, handling 424
statistics
  client for AIX 347
  client for Windows 347
statserver 286
status
  cluster 278
  engine 282
  Metadata server 286
status lights 366
stfsclient 303
stfsdebug 345
stfsdisk 306
stfsdriver 309
stfsmount 311
stfsstat 347
stfsstatus 313
stfsumount 314
stopautorestart 289
stopCimom 351
stopcluster 290
stopConsole 352
stopengine 291
stopmetadatacheck 292
stopping
  all Metadata servers in the cluster 290
  metadata checker 292
  SAN File System console Web server 327, 352
stopping an Metadata server 293
stopping the Administrative agent 351
stopserver 293
storage engine, replacing components 411
storage pool
  adding volumes to 240
  creating 238
  deleting 256
  listing 216
  setting the default 262
support
  general 151
  online 152
  telephone 152
suspending a volume 294
suspendvol 294
switch
  block 369
  power-on password override 369
symptoms
  beep 377
System error LED 365
system planar
  connectors
    user-installable options 371, 372
  LEDs 369
  switch block 369
system planar and shuttle, replacing 142
system reliability 425
System-identification LED 365
system-planar
  connectors
    external port 367
    internal cable 368
    SCSI 368, 374

T

Taiwan electrical emission statement 417
tank extractbootrecord 353
tank lscluster 355
tank lsdisklabel 358
tank lsversion 357
tank resetcluster 360
tank resetversion 362
Tank.Bootstrap file 353, 360
tanktool
  ending the session 185, 244
  getting help 186
temperature
  error messages 399
  shutdown messages 400
thermal material 131
TRACE command 328
tracing
  Metadata server 40
trademarks 427
traps
  listing settings 226
troubleshooting
  Administrative server 67
  CD-ROM 88
  client performance 80
  diskette drive 89
  Ethernet controller 90
  hard disk drive 92
  keyboard 93
  local network 59
  LUNs 77
  memory 94
  Metadata server 56

troubleshooting *(continued)*
microprocessor 94
monitor 95
mouse 93
power 97
power-off 100
power-on 100
Universal Serial Bus (USB) 99
user access, command-line interface 70
user access, commands 71
user access, console 68
truststore
creating 335

U

United Kingdom telecommunications safety requirement 415
United States electronic emission Class A notice 414
United States FCC Class A notice 414
Universal Serial Bus (USB) port 375
UNIX-based clients 12
unmounting the global namespace on a client for AIX 314
upgrading cluster 295
usepolicy 296
utility
startCimom 343
stopCimom 351

V

video port 375
vital product data (VPD)
hardware 46
software 47
voltage regulator module (VRM), installing 131
voltage-related shutdown 402
volume
activating 163
adding to a storage pool 240
listing files in 245
modifying settings 178
removing 258
scanning SAN File System for 306
suspending 294
VPD (vital product data)
hardware 46
software 47
vpdlog command 35

W

Web site
compatible options 426
SCSI standards 375
Web sites x, 151
Who should use this guide viii
Windows client
command 315
Windows-based client
administrative privileges 14
case sensitivity 14
Windows-based client *(continued)*
command line interface (CLI) 15
file management 15
NTFS features 15
Windows-based clients 13
Windows, client for
statistics 347
working inside engine with power on 426

Readers' Comments — We'd Like to Hear from You

IBM TotalStorage SAN File System
(based on IBM Storage Tank™ technology)
Maintenance and Problem Determination Guide
Version 1 Release 1

Publication No. GA27-4318-00

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	☐	☐	☐	☐	☐

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	☐	☐	☐	☐	☐
Complete	☐	☐	☐	☐	☐
Easy to find	☐	☐	☐	☐	☐
Easy to understand	☐	☐	☐	☐	☐
Well organized	☐	☐	☐	☐	☐
Applicable to your tasks	☐	☐	☐	☐	☐

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? ☐ Yes ☐ No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Name _____ Address _____

Company or Organization _____

Phone No. _____

Readers' Comments — We'd Like to Hear from You
GA27-4318-00

Cut or Fold
Along Line

Fold and Tape | **Please do not staple** | Fold and Tape

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM Corp
Dept. CGFA
PO Box 12195
Research Triangle Park, NC 27709-9990

Fold and Tape | **Please do not staple** | Fold and Tape

GA27-4318-00

Cut or Fold
Along Line

Printed in U.S.A.

GA27-4318-00